The Sentiment of Reserve Bank of New Zealand Monetary Policy Media Releases and the Impact on Financial Markets

Kate Allen

2025

A dissertation submitted to Auckland University of Technology in partial fulfilment of the requirement for the degree of Master of Business (MBus).

Supervisors: Aaron Gilbert and Stephanie Rossouw

Abstract

Central banks play a critical role in shaping financial markets through their monetary policy decisions. As part of their commitment to transparency, they provide forward guidance through, for example, soft information such as Official Cash Rate announcements, which convey the tone and sentiment of the announcement. Tone and sentiment influence market participants' expectations and, in turn, market behaviour. Therefore, understanding how sentiment impacts financial indicators is vital for both policymakers and investors. This study investigates whether the sentiment of the Reserve Bank of New Zealand's Monetary Policy (MP) media releases influences financial markets, specifically the stock market and exchange rate. Sentiment analysis is performed using three methods: the Valence Aware Dictionary for sEntiment Reasoner (VADER) and two bag-of-words lexicons—Loughran-McDonald (LM) and Hu and Liu (HL). Sentiment variables are incorporated into an event-study framework using OLS regression analysis to identify associations between sentiment and financial outcomes. The results showed that VADER slightly outperforms LM and HL. Results from the event study suggest that a marginally negative association exists between the VADER negative score and the New Zealand stock market adjusted for the US stock market. Therefore, a higher proportion of negative sentiment in an MP media release could negatively impact the New Zealand stock market. Similarly, there was a stronger negative association between the VADER positive score and the trade-weighted index, indicating that an increase in the proportion of positive sentiment in an MP media release depreciates the New Zealand exchange rate. This indicates that positive sentiment might lead to increased domestic confidence and spending, potentially raising imports. This contribution provides valuable insights into central bank communication strategies in small, open economies.

Acknowledgements

Completing this dissertation has been a challenging yet rewarding journey, and I am sincerely grateful to those who supported me along the way.

Firstly, I would like to thank my supervisors, Aaron Gilbert and Stephanie Rossouw, for their willingness to oversee a dissertation on a topic of my own choosing. Their guidance, expertise, and patience were instrumental throughout this process, and I truly appreciate the time and effort they dedicated to helping me complete this work.

I am also deeply thankful to my family—my parents, Steve and Brenda Allen, and my sister, Lucy Allen—for their steady encouragement throughout my master's program.

Special thanks go to my former colleague, Stephanie Mitchell, for her help with coding and for introducing me to Python, which was essential for this research.

I am also grateful to my colleagues and managers at the Reserve Bank of New Zealand for their understanding and support, allowing me to balance the demands of work, health, and completing this dissertation.

Lastly, I would like to thank everyone who helped and encouraged me throughout this process.

Thank you.

Table of Contents

List of Tables

List of Figures

Attestation of Authorship

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person (except where explicitly defined in the acknowledgements) nor used artificial intelligence tools or generative artificial intelligence tools (unless it is clearly stated, and referenced, along with the purpose of use), nor material which to a substantial extent has been submitted for the award of any other degree or diploma of a university or other institution of higher learning.

Students Signature

Date: 20/01/2025

Student ID: 21146693

Chapter 1 Introduction

Central banks play a critical role in shaping financial markets through their monetary policy decisions. Interest rates directly impact financial markets, with higher interest rates often leading to stock price declines and exchange rate appreciations. To minimise uncertainty, central banks aim to be transparent and even slightly predictable when setting interest rates. A part of their commitment to transparency is providing forward guidance on their future actions based on current economic conditions. In the process of providing forward guidance, there are two sets of information being provided: so-called hard information, which represents quantifiable information, and soft information, which, in the case of Official Cash Rate (OCR) announcements (Monetary Policy media releases) is the tone and sentiment conveyed in the announcement. Tone and sentiment are closely studied and reported on and influence market participants' expectations and, in turn, market behaviour. Therefore, understanding how sentiment impacts financial indicators is vital for both policymakers and investors.

This study investigates whether the sentiment of the Reserve Bank of New Zealand’s (RBNZ) Monetary Policy (MP) media releases influences financial markets, specifically the stock market and exchange rate. This study focuses on the stock market and exchange rate as proxies to measure the impact of sentiment on financial markets due to their sensitivity to monetary policy announcements. This study justifies the choice of using the stock market because of how it reflects investor confidence and expectations about economic conditions (Zeng et al., 2025), while the exchange rate is directly impacted by changes in interest rates through capital flows (Shahzad et al., 2020). Therefore, both the stock market and exchange rate are key metrics for evaluating the transmission of monetary policy in New Zealand.

Both the reaction of stock markets (Gorodnichenko et al., 2021; Parle, 2022; Anderson and Thorshaug, 2023) and exchange rates (Gorodnichenko et al., 2021; Dowey, 2024) have been studied previously in other countries within the context of sentiment analysis, monetary policy documentation, and financial markets. By focusing on the stock market and exchange rate, this study aligns with existing research while providing new insights specific to New Zealand.

In this study, the sentiment analysis is performed using three methods: the Valence Aware Dictionary for sEntiment Reasoner (VADER) (Hutto & Gilbert, 2014) and two bag-of-words lexicons—Loughran-McDonald (LM) (Loughran & McDonald, 2011) and Hu and Liu (HL) (Hu & Liu, 2004). Sentiment variables are incorporated into an event-study framework using OLS regression analysis to identify associations between sentiment and financial outcomes. Refer to section 4 for detailed information on the methodology employed.

Previous studies (detailed in the literature review in section 2) that examined the sentiment of central bank communication and its impacts on financial markets have generated mixed results. Parle (2022) found that a hawkish tone in European Central Bank press conferences positively impacted stock prices, suggesting optimism about future economic conditions was viewed as leading to better future company performance. Conversely, Gorodnichenko et al. (2022) found economically significant but no statistically significant effects of sentiment on stock markets, suggesting a positive relationship with a dovish tone. Other studies, such as Anderson and Thorshaug (2023), identified links between negative sentiment and higher stock market returns, while Dowey (2024) and Gorodnichenko et al. (2021) highlighted sentiment's impact on exchange rates, including appreciations and depreciations under varying conditions.

To date, there have been no New Zealand-based studies that examine the sentiment of central bank communication and its impact on financial markets. Cherry and Tong (2023) performed a sentiment analysis study on New Zealand monetary policy statements using three bag-of-words methodologies, including the LM and HL lexicons. In this study, I aim to expand on Cherry and Tong's (2023) study by including VADER, a more accurate measure of sentiment that does not suffer from the weaknesses associated with bag-of-words lexicons, such as the LM and HL methods. The aforementioned will allow me to determine whether the sentiment of the monetary policy communication has an impact on New Zealand's financial markets, which has not been studied in the New Zealand context until this paper.

The sentiment results indicate that the Monetary Policy (MP) media releases from the RBNZ were predominately neutral when measured across the three sentiment measures (VADER, LM, and HL). Even though the results were mostly neutral, VADER produced results with more positive sentiment than negative sentiment, with a high compound score in the positive range on average and a positive sentiment score higher than the negative sentiment score on average. This indicates that the MP media releases had a higher proportion of positive sentiment than negative sentiment on average. In contrast, the LM bag-of-words indicated net negative sentiment scores and ratios, and the HL method showed results that were very close to zero, even though slightly more positive.

The results from the event study regression analysis showed that VADER sentiment variables provided more reliable results and demonstrated that there is an association between the monetary policy media releases and financial indicators. There was a marginally negative association between the VADER negative score and the New Zealand stock market adjusted for the US stock market (NZX SPY). This result indicates that a higher proportion of negative sentiment in a monetary policy media release could negatively influence the New Zealand stock market. This result aligns with the expectation that pessimistic media releases may dampen investor confidence. There was a stronger negative association between the VADER positive score and the trade-weighted index (TWI), indicating that an increase in the proportion of positive sentiment in an MP media release could depreciate the New Zealand exchange rate (TWI). This stronger negative association suggests that optimistic monetary policy releases might signal to markets that the RBNZ is adopting a dovish stance (e.g., signalling low rates or economic challenges ahead), which could reduce capital inflows, depreciating the New Zealand Dollar.

Results from using the LM sentiment variables indicate that an increase in net positive sentiment could lead to an increase in the TWI returns on average, which contradicts the

VADER results. However, this contradiction with VADER may arise from LM's lack of contextual understanding (e.g., "positive" words in dovish contexts not being truly market-positive).

Given the results, this study contributes to the literature in several ways. First, it expands on existing research by applying VADER (in addition to the LM and HL lexicons) to New Zealand monetary policy communication (MP media releases), which has not been explored before. This allowed me to test whether utilising VADER will provide more accurate sentiment results. Second, the study adds to the limited literature on sentiment analysis in a New Zealand context, incorporating a regression analysis to explore relationships between sentiment and financial indicators. Third, this study contributes to the literature by providing evidence regarding the impact of sentiment analysis on financial markets, which is currently showing mixed results. This contribution provides valuable insights into central bank communication strategies in small, open economies.

The remainder of the paper is organised as follows. Section 2 provides a literature review on natural language processing, sentiment analysis, and its applications to central bank communication and financial markets. The data and selected variables are discussed in section 3. The methodology is outlined in section 4, followed by the results and analysis in section 5. The paper concludes with section 6.

Chapter 2 Literature Review

2.1 Introduction

Central banks can impact markets when they change their interest rates. However, they aim to mitigate any significant movements by being forthcoming with their forecasts and providing forward guidance on their future intentions. As a result, investors and markets are very interested in the actual decisions made, central banks' thinking, and the so-called soft information derived from the sentiment of announcements. Therefore, the way that central banks articulate information regarding monetary policy, specifically the sentiment conveyed in this communication, is important. A growing body of literature has investigated how the words used to convey information impact how people respond to information, including the response of financial markets.

In this section, I synthesise a body of literature that focuses on investigating whether the sentiment of monetary policy communication or associated documentation impacts financial and economic indicators. The literature review begins by introducing Natural Language Processing (NLP) and focuses on different lexicons used for sentiment analysis. Second, I discuss literature that established previous relationships between sentiment and financial indicators. This section concludes by highlighting limitations from previous studies that this study aims to address.

2.2 Natural Language Processing

Natural language processing (NLP) is the methodology adopted to perform text sentiment analysis. Sentiment analysis is an automated process of determining whether a text expresses a positive, negative or neutral opinion about a topic (Hailong et al., 2014). Rossouw and Greyling (2021) categorise natural language processing into three different methods:

(I) Rules-based systems: using, for example, lexicons.

(II) Automatic systems: relying on machine learning techniques to learn from data.

(III) Hybrid systems: combining both rules-based and automatic approaches.

This study uses rules-based and automatic systems and will focus subsequent discussions accordingly.

2.2.1 Rules-based systems – Lexicons

Within rules-based systems, there are three subcomponents which are outlined by Hailong et al. (2014):

(I) Manual approach

(II) Dictionary-based approach

(III) Corpus-based approach

The first approach involves someone manually defining what they want to achieve with the NLP. They then create a set of rules for each word or particular word patterns and structures, process the text via tokenisation, apply the rules to sample texts and refine the rules, and then finally apply the rules and methods to the actual text they want to analyse the sentiment. The manual approach is very time-consuming and, therefore, is not usually used alone. Instead, it is combined with automated approaches as a final check to address any mistakes made by automated methods (Hailong et al., 2014). In automated methods, the results depend on the domain and the quantity and quality of the training dataset. Some machine learning methods struggle with context-specific sentiment. While a human can manually make sure that the rules are applied correctly to the text, this is time-consuming.

In contrast, the corpus-based methods use a seed set of sentiment words, an initial set of words with known polarity and exploit syntactic patterns of co-occurrence to identify new sentiment words and their polarity in a large corpus (Hailong et al., 2014). This approach is not as popular as the dictionary-based approach because it requires a large body of text to identify all the potential additional sentiment words and phrases. Even though there is a given seed set in the corpus-based approach, this is just the starting point. The corpus needs to be large enough to cover all the possible combinations of text to have reliable sentiment analysis. Based on the disadvantages discussed above, the manual or corpus-based rules-based sentiment analysis will not be utilised in this study.

The dictionary-based approach is widely used in the text sentiment literature; see, for example, Almeida et al. (2021), Cherry and Tong (2023) and Nguyen and Huynh (2020). The main strategy in this approach is to collect an initial seed set of sentiment words and their orientation manually. Then, these are searched for in a dictionary to find their synonyms and antonyms and expand the original seed set (Hailong et al. 2014). A key advantage is that a seed set of words has already been identified and, as such, can be directly applied to a text to determine its sentiment. Some examples of the dictionary-based approach seed sets (referred to as lexicons) include Loughran and McDonald (LM) (2011), Hu and Liu (HL) (2004), the Harvard General Inquirer (GI), and the Valence Aware Dictionary for sEntiment Reasoner (VADER) (Hutto & Gilbert, 2014). The LM and HL lexicons are both widely used within business contexts. The LM lexicon was developed specifically for the field of economics and finance. The HL lexicon was developed after gathering customer reviews of products sold online.

The LM and HL are both lexicons widely used in conjunction with the so-called "bag-of-words" method. The "bag-of-words" method is often used in sentiment analysis and is one of the more straightforward methods which treats a document simply as a collection of words. The lexicon has a dictionary of words, where each word is classified as having a positive or negative sentiment. The individual words' sentiment is then accumulated over the entire text to get the overall sentiment. For instance, a document with more positive words than negative words would be deemed to have a positive sentiment. These "bag-of-word" lexicons have been used in several studies, including, for example, Cherry and Tong (2023), Segawa (2021), Shapiro et al. (2019), and Shapiro and Wilson (2021). There are several disadvantages of the bag-of-words approach, including that words are often equally weighted, meaning positive words like good and great may be given the same weight even though great would generally be considered to convey a stronger sentiment. Additionally, the bag-of-words typically look at each word in isolation of its context or place in a sentence's meaning. For example, “I like something” and “I do not like something” would have similar scores.

A more sophisticated lexicon, rules-based sentiment analysis tool is VADER. VADER which consists of both a lexicon—a list of several thousand words ("unigrams") labelled from -4 to 4 corresponding to a most negative to a most positive score— and a set of heuristic rules that account for a word's context within the sentence. VADER assigns a (net) negativity score to a sentence by aggregating negativity scores of words within the sentence (Shapiro et al., 2019; Shapiro & Wilson, 2021).

Lexicon-based sentiment analysis has been practically applied in a range of fields, including politics, marketing, subjective well-being, finance and economics. In politics, Abercrombie and Batista-Navarro (2020) compiled a systematic literature review of 61 studies on sentiment analysis of parliamentary debates. They were able to use the sentiment and the position of the political party to determine the future sentiment of parliamentary debates. Hu and Liu (2004) extrapolated customer reviews for electronic devices sold online and manually created a lexicon known as the Hu and Liu (HL) lexicon to use for sentiment analysis targeting future marketing strategies. Rossouw and Greyling (2020) used a combination of lexicons, including NRC (National Research Council of Canada Emotion Lexicon developed by Turney and Mohammad (2010)), VADER, TextBlob, Syuzhet, AFINN, and Bing to analyse the emotion and the sentiment of social media posts on Twitter, effectively measuring the mood of a nation. Twitter has also been utilised to predict the stock market in multiple developed and developing countries, which was performed by Greyling et al. (2019); other studies involving VADER sentiment analysis and the stock market include Cristescu et al. (2022), Das et al. (2022), Deveikyte et al. (2022), Li et al. (2020) and Soni and Mathur (2023). Additional studies that focused on finance and economics utilising sentiment analysis include Anderson and Thorshaug (2023), Cherry and Tong (2023), Dowey (2024), Gorodnichenko et al. (2021), Parle (2022), Segawa (2023) and Shapiro and Wilson (2021). These studies will be discussed in more detail in section 2.3.

2.3 Sentiment Analysis of Monetary Policy Documentation

There is a growing body of literature in economics which applies sentiment analysis to monetary policy documentation and communications. Several of these studies have been conducted in the USA, where sentiment analysis was performed on monetary documents from the Federal Reserve. Other studies focused on Europe, and one focused on New Zealand. As is seen in the discussion below, using the sentiment of central bank communication to determine its effect on financial markets and economic indicators has generated mixed results.

Gorodnichenko et al. (2021) used BERT (Devlin et al.,2019) to analyse Federal Open Market Committee (FOMC) statements and transcripts from Q and A sessions. Using a dove-hawk-styled[1] index, Gorodnichenko et al. (2021) were able to determine whether a monetary policy document is more dovish or hawkish. Gorodnichenko et al. (2021) performed panel regression analysis and found an economically significant but statistically not significant positive relationship between dovish sentiment and stock market returns. In terms of the relationship between sentiment and exchange rates, Gorodnichenko et al. (2021) found that a more dovish tone appreciated the US Dollar against the Euro, but there was not much response against the yen. As for the relationship between dovish tone and inflation expectations, Gorodnichenko et al. (2021) found a positive but not statistically significant relationship.

Similarly, Parle (2022) adopted a dynamic topic model (DTM) to perform sentiment analysis and categorised the text of European Central Bank press conference transcripts into a hawk-dove-styled index. Parle (2022) adopted an event study and found that a more hawkish tone at ECB press conferences is associated with a positive impact on

[1] In the context of central banking and monetary policy, the terms "hawkish" and "doveish" are often used to describe the stance or outlook of policymakers regarding interest rates and economic conditions, A "hawkish" stance usually indicates a stricter or more aggressive approach to monetary policy and is often associated with monetary tightening policy (increase in rates). A "doveish" stance reflects a more cautious approach to monetary policy and is often associated with loosening monetary policy (decrease in rates). A dove-hawk-styled index categorises words as either being hawkish or doveish. In the case of Gorodnichenko et al. (2021), the doveish words are given a positive score (1 to 10), and the hawkish words are given a negative score (-1 to -10).

stock prices, which contradicts Gorodnichenko et al.'s (2021) findings. This could be due to the way they constructed their dove-hawk indexes, the different language used across the two central banks and because Gorodnichenko et al. (2021) had a large portion of their study focused on voice tone (i.e. noise waves). As regards stock prices, Parle (2022), using an event study, found that a more hawkish tone at ECB press conferences is associated with a positive impact, contradicting Gorodnichenko et al.'s (2021) findings. Additionally, Parle (2022) found correlations between the sentiment indexes created and the Gross Domestic Product (GDP), as well as inflation, suggesting these could be used to predict GDP and inflation.

Dowey (2024) adopted a large language model to analyse the sentiment of Riksbank monetary policy reports. Similar to Gorodnichenko et al. (2021) and Parle (2022), Dowey (2024) created a classification where the words are either hawkish or dovish, as well as creating an index for positive, negative, and neutral sentiment. Dowey (2024) performed an event study and found that positive monetary policy sentiment significantly increased treasury bills (short-maturity debt) in Sweden. However, Gorodnichenko et al. (2021) and Parle (2022) did not find statistically or economically significant results between the sentiment and bonds (long-term maturity debt). Investigating the relationship between the sentiment of monetary policy documentation and exchange rates, Dowey (2024) found mixed results when comparing positive and negative sentiment to changes in the Swedish Krona against the Euro and the US Dollar. For the positive sentiment, Dowey (2024) found a slight downward trend for the SEK/USD and SEK/EUR, although the results were not statistically significant. However, they were able to determine a statistically significant relationship between negative sentiment and increases in the Swedish Krona against the USD and EUR.

Anderson and Thorshaug (2023) adopted Multinominal Inverse Regression (MNIR) and Naive Bayes classifiers to create their own dictionaries. They applied the "bag-of-words" approach to obtain the sentiment of the Norges Bank Monetary Policy Evaluations (MPEs). Their results indicate that negative sentiment leads to higher stock returns.

Segawa (2023) adopted the LM dictionary and utilised TextBlob to perform sentiment analysis on the Monetary Policy Committee (MPC) statements from the South African Reserve Bank (SARB). Adopting a vector autoregression (VAR) model, Segawa (2023) determined a causal relationship between the tone and sentiment of the South African Reserve Bank (SARB) and inflation expectations.

Additionally, Shapiro and Wilson (2021) employed a regression model that regressed the sentiment of FOMC transcripts as the dependent variable (a proxy for a central bank loss) against independent variables, including the inflation gap, unemployment gap, output growth and stock market performance. The sentiment measures used were the LM method and VADER. The authors were able to quantify how the changes in sentiment expressed by FOMC members were related to the Federal Reserve's loss function and preferences. Overall, Shapiro and Wilson (2021) found that the FOMC had an implicit inflation target averaging 1.5%, which is below the 2% inflation target and that the central bank's loss function was influenced more by output growth and stock market performance than the perception of current economic slack (unemployment, their second objective function).

In New Zealand, Cherry and Tong (2023) performed sentiment analysis on monetary policy statements (MPS) using the LM, HL and GI dictionaries. Cherry and Tong (2023) found using the three lexicons that the Reserve Bank of New Zealand’s (RBNZ) monetary policy statements had predominately neutral tone and that less than 6% of each MPS carried sentiment. In this study, the Cherry and Tong (2023) sentiment results will be used as a benchmark for the sentiment analysis on monetary policy media releases from the RBNZ whilst expanding on the study by utilising VADER and performing regression analysis to determine whether there is an association between the sentiment in the Monetary Policy (MP) media releases and financial markets (New Zealand stock market and exchange rate).

VADER adds value to sentiment analysis as it does not look at individual words in isolation like the other LM, HL and GI dictionaries but accounts for the structure and order of the words in full sentences. Using the LM method for its financial-based dictionary in conjunction with VADER, which accounts for the overall text, is a robust way to determine the sentiment of economic or financial text. VADER is not perfect since its lexicons are created from social media text, but it can be effective with the LM method. This will align my study with that of Shapiro and Wilson (2021), who is the only study that adopted VADER, as well as the LM lexicon, in the context of sentiment analysis of monetary policy documents.

2.4 Summary

There is a range of literature on monetary policy sentiment analysis across the United States of America, Europe, South Africa, and New Zealand. The studies that adopted lexicon-based methods are vastly different in terms of the financial regression methodology used. Cherry and Tong (2023) used only sentiment analysis and did not conduct regression analysis focusing on financial indicators. Segawa (2023) focused on inflation expectations using a VAR model with their LM sentiment scores. Shapiro and Wilson (2021), however, are one of the only monetary policy studies that incorporate VADER in their sentiment analysis.

Therefore, a limitation of previous studies that this study aims to address is utilising the Loughran and McDonald (2011) lexicon as well as VADER to determine which provides the most reliable information regarding the sentiment of monetary policy documentation. Second, the sentiment scores will be used as a regressor (independent variable) in the regression analysis, specifically an event-styled regression analysis, something Cherry and Tong (2023) did not do. The analysis in this study contributes to New Zealand's understanding of the role that sentiment plays in predicting financial factors. Therefore, this study extends the work done by Cherry and Tong (2023) by including VADER and performing an event study to determine whether an association between sentiment, the stock market and the exchange rate exists.

Chapter 3 Data and variables

This section outlines the data sources, the time frame under investigation and the variables constructed to address the question of whether an association exists between the sentiment of monetary policy media releases and financial indicators. The first section details the sentiment data and sentiment covariates, and the second section details the financial and economic data, including the dependent variables and the control covariates.

3.1 Constructing Sentiment Analysis Data from Monetary Policy Media Release

The Reserve Bank of New Zealand (RBNZ) reviews the official cash rate (OCR) seven times a year, roughly every six weeks. The OCR announcement is made at 2 pm on a Wednesday and is made via a Monetary Policy (MP) media release. In four of the seven reviews, once per quarter, the MP media release is accompanied by a Monetary Policy Statement (MPS) and a televised press conference. The monetary policy statement is a substantial document that provides considerable detail on the RBNZ’s view of the country's economic state.

Initially, 136 monetary policy media releases were manually obtained from the RBNZ official website, spanning the period 26 January 2006 to 16 August 2023. This date range was initially chosen since it included all MP media releases available on the website, even though the OCR was introduced on 17 March 1999. This study commenced in July 2023, so the latest date chosen was 16 August 2023 to include the most up-to-date announcements.

The MP media releases have evolved over time, and the process of how they are released has also changed. In the earlier media releases, they include direct quotes from the governor’s speeches that he makes at the televised press conferences. In more recent years, the media release has been released 1 hour before the televised press conference and has a more formal format. There is a different monetary policy secretary who constructs this media release each round.

In 2019, the RBNZ introduced a monetary policy committee (MPC). From 8 May 2019 onwards, the decision regarding the OCR is made by the MPC instead of the governor alone. The MP media releases also changed during this time. The first half of the media release follows the previous format on the decision and the state of the economy. However, there is a second half where there is a summary record of the meeting, which is essentially the minutes from the MPC meeting. Different people also write these two halves of one MP media release. Subsequently, the date range was refined from 26 January 2006 to 27 March 2019 instead of 16 August 2023 due to this dramatic change in format, which could potentially skew the sentiment results. The date of 27 March 2019 represents the final OCR announcement in the form of an MP media release before the MPC was introduced. Another reason why those later dates were excluded was the COVID-19 pandemic, which brought about the introduction of the large-scale asset purchase plan (LSAP). The LSAP was a plan by the RBNZ to purchase up to $100 billion worth of government and other bonds. The LSAP program was halted on 23 July 2021, and the RBNZ has since been gradually selling these bonds. Based on the combination of the introduction of the MPC, the summary record of meetings, COVID-19 and the LSAPs, the final date range chosen for the Monetary Policy (MP) media releases is from 26 January 2006 until 27 March 2019, which includes 101 MP media releases.

The MP media releases were copied into .txt files and cleaned before using Python to extract the sentiment; refer to the next section below for the construction of the sentiment variables. For the initial cleaning process, the following had to be removed: the title, the date, the attendees, the media contact information, and I translated or removed any use of Te Reo Māori (as the lexicons would not be able to interpret this). Extra dates, titles and names are not useful in the sentiment analysis process; therefore, they were removed.

3.1.1 Sentiment covariates

To determine whether the sentiment of MP media releases is associated with the stock market and or exchange rate, I needed to construct the sentiment covariates. Ten sentiment covariates are constructed and divided into three groups depending on the methodology used. In this study, the primary measures of sentiment are the VADER scores, while also looking at the Loughran McDonald (LM) ratios and Hu and Liu (HL) ratios, which have been used more extensively but fail to consider the context of words when calculating sentiment.

3.1.1.1 VADER covariates

One of the contributions of this study is the use of VADER to calculate sentiment in the MP media releases. VADER comes with a number of significant benefits over bag-of-words lexicons, such as those applied by Cherry and Tong (2023), Dowey (2024), Gorodnichenko et al. (2021), Parle (2022), and Anderson and Thorshaug (2023). VADER uses an existing lexicon that considers both whether a word is positive or negative and the relative strength of the word, as well as the context of words, to generate a more accurate measure of sentiment. Additionally, VADER does not have the issue of finding a large corpus of documents to train the sentiment model, allowing us to avoid losing media releases to train the model. Greyling et al. (2019) successfully used VADER sentiment of tweets from Twitter to predict the stock market, and Shapiro and Wilson (2021) utilised VADER on federal reserve documentation to generate sentiment measures that were used to estimate a central bank loss function. To the best of my knowledge, no study exists that specifically used VADER to calculate the sentiment of reserve bank communications to examine the impact on financial variables, such as the stock market and exchange rate.

VADER (Valence Aware Dictionary and sEntiment Reasoner) consists of both a lexicon—a list of several thousand words ("unigrams") labelled from -4 to 4 corresponding to most negative to most positive— and a set of heuristic rules that account for a word's context within the sentence. VADER assigns a (net) negativity score to a sentence by aggregating across negativity scores of words within the sentence (Shapiro et al., 2019). In addition to this, VADER changes the score based on five rules: capitalisation, negation, punctuation, being proceeded by versus following the word “but” and whether a word is proceeded by a degree modifier such as “very”, “extremely”, “slightly” etc. (Shapiro et al., 2019).

VADER provides four scores: positive, negative, neutral and a compound score. These four scores are automatically generated through the estimation code and are subsequently used as the sentiment covariates.

The VADER positive, negative and neutral scores sum to 1. Each score gives you the proportion of the MP media release that has positive, negative or neutral sentiment. The remaining VADER score is the compound score, which is calculated by normalising the sum of valence scores (-4 to 4) from individual words. The result of these scores being aggregated and normalised is a compound score between -1 and 1 (Akladyous, 2023). A compound score is positive if it is greater than 0.05, neutral if it’s between -0.05 and 0.05 and negative if it is less than -0.05.

VADER Compound Score	$compound = \frac{sum_s}{\sqrt{sum_s} + \alpha}$
Positive Compound Score	$X \geq 0.05$
Neutral Compound Score	$-0.05 < X < 0.05$
Negative Compound Score	$X \leq -0.05$

Refer to Table 1 for the VADER descriptive statistics.

Table 1: VADER Descriptive Statistics

Descriptive Statistics	**N**	**Mean**	**Median**	**SD**	**Min**	**Max**
VADER Compound Score	101	0.574	0.920	0.607	-0.910	0.994
VADER Positive Score	101	0.107	0.107	0.032	0.042	0.189
VADER Negative Score	101	0.070	0.067	0.023	0.000	0.130
VADER Neutral Score	101	0.823	0.820	0.037	0.755	0.947
Word Count						
VADER Word Count	101	296	303	60	144	433

Source: Author's calculations

Table 1 shows that the VADER compound score is positive on average, with a mean of 0.57, larger than the upper bound of 0.05 for a document to be considered neutral (-0.05<neutral<0.05). The median score compound score is positive, with a score of 0.92. Even though the compound score is 0.57 and the median is 0.92, there is a large range in the sentiment scores for individual Monetary Policy (MP) media releases ranging from -0.91 (which is close to the most negative end of the compound scale of -1) and 0.994 (which is nearly entirely positive – close to 1). The minimum of -0.91 was for the MP media release made on 30 April 2009, approximately 7 months after the global financial crisis (GFC) started with the collapse of the Lehman Brothers. This media release did not include a televised press conference, and the official cash rate (OCR) was reduced by 50 basis points. Dr Allan Bollard mentions in the media release that due to the scale of the GFC and domestic circumstances, it would likely be some time before economic activity returns to robust and healthy levels. The maximum compound score of 0.994 was also in 2009, but this media release was on 10 December 2009 and was supplemented with a televised press conference. Dr Bollard mentioned in his speech that business confidence is improving and that both fiscal and monetary policy support is assisting the economy. There was no change in the OCR on 10 December 2009.

The VADER Positive Score represents the proportion of text that has positive sentiment. In this case, the mean and median scores are the same at 10.7%. The proportion of positive sentiment in Monetary Policy media releases ranges from 4.2% to 18.9%. The minimum of 4.2% was the 24 July 2008 MP media release during the GFC and the maximum on 10 December 2009, which was also the date where the compound score maximum occurred (see Figure 1).

The VADER Negative Score represents the proportion of the text that has negative sentiment. On average, the negative proportion of each Monetary Policy media release is almost 7%; this ranges between 0% of the text having negative sentiment and a maximum of 13% of the text being negative. The minimum negative score fell on 26 April

2012, and the maximum on 10 November 2016. The media release on 26 April 2012 was only a monetary policy review (no televised press conference or MPS), and the OCR remained unchanged. On 10 November 2016, the OCR was reduced by 25 basis points due to weak inflation and emphasis on economic uncertainty during this period (refer to Figure 1). This media release was accompanied by a Monetary Policy Statement (MPS) and a televised press conference.

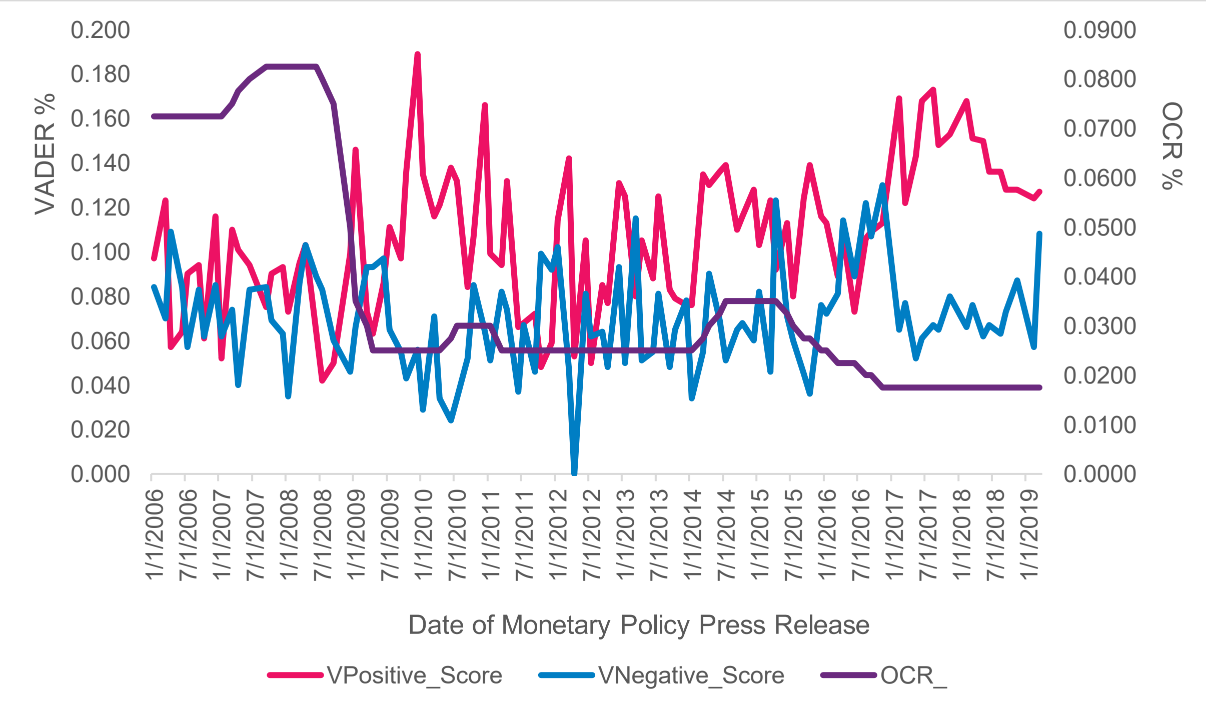

Figure 1: VADER positive and negative scores and the OCR from 2006-2019

Source: Author's calculations and RBNZ OCR data

Table 1 shows the VADER Neutral Score, which shows that the average proportion of neutrality is 82%. The range of neutral sentiment of monetary policy media releases is between almost 66% and 95%. This result is consistent with Cherry and Tong's (2023) finding that the Reserve Bank of New Zealand's (RBNZ) monetary policy statements (MPS) are very neutral when the sentiment is determined using the bag-of-words method.

What is interesting about these VADER scores is that they cannot be interpreted in isolation. The compound score shows that the sentiment of the MP media releases is positive on average. However, the neutral score tells us that the Monetary Policy media releases are predominantly neutral on average (82% of the text).

In the later regression models, this study uses the three VADER covariates: the compound, positive, and negative scores. This study considers both the impact of the VADER compound score and the impact of the VADER positive and negative scores separately but jointly.

3.1.1.2 Bag-of-words covariates

The "bag-of-words" method is frequently used in sentiment analysis and is one of the more straightforward methods. It is based on a dictionary of words where each word is classified as having a positive or negative sentiment. The individual words' sentiment is then accumulated over the entire text to get the overall sentiment. The bag-of-words methodology does not account for the sentence structure or how the words are ordered, relying solely on various ratios of the number of positive to negative words to determine sentiment.

To prepare the MP media releases for analysis, they had to be “tokenise,” meaning the text was broken down into individual words, punctuation, numbers, and names (GeeksforGeeks, 2024). Python’s NLTK (Natural Language Toolkit) was utilised to tokenise the text and remove stopwords. Stopwords include the following types of words: articles, prepositions, conjunctions, pronouns and auxiliary words. Some examples of stopwords include: a, the, in, at, and, but, or, I, you, she, is, am and was. NLTK removed punctuation, special characters, numbers, names and white spaces to produce clean text. Once the text is “clean”, one is left with a “bag-of-words” ready for sentiment analysis.

To perform sentiment analysis using the bag-of-words methodology, you must adopt a lexicon, which is a predetermined dictionary of words. In the lexicons, the words are categorised as either positive, negative, or neutral. In this study, two lexicons were adopted: Loughran McDonald (LM) and Hu and Liu (HL). Four scores/ratios are calculated: the sentiment score, sentiment ratio, subjectivity ratio and polarity ratio.

The sentiment score is all the positive words minus all the negative words. There are also neutral words included in the overall text, but neutral words do not affect this calculation.

The sentiment ratio is the number of positive words minus the number of negative words divided by the number of total words (cleaned, tokenised and including neutral words). The result is the net positive sentiment as a proportion of the overall text.

$$Sentiment\ ratio_t = \frac{Positive_t - Negative_t}{Nt} \tag{1}$$

The subjectivity ratio is the number of positive words plus the number of negative words (total words with sentiment) as a percentage of the total number of words in the text (cleaned, tokenised and including neutral words).

$$Subjectivity\ ratio_t = \left(\frac{Positive_t + Negative_t}{Nt}\right) * 100 \tag{2}$$

The polarity ratio is the net positive words over the number of words with sentiment. The polarity ratio falls between the range of -1 to 1. A score from -1 to -0.5 usually indicates negative sentiment, a score from -0.5 to 0.5 usually indicates neutral sentiment, and a score from 0.5 to 1 usually indicates positive sentiment.

$$Polarity\ ratio_t = \frac{Positive_t - Negative_t}{Positive_t + Negative_t} \tag{3}$$

3.1.1.2.1 Loughran McDonald (LM) covariates

In the first bag-of-words Python script, the Loughran McDonald (LM) lexicon was applied to conduct sentiment analysis. Loughran and McDonald (2011) built their lexicon based on 10-K filings from 1994 until 2008. This lexicon has a financial and economic focus. They developed six different word lists: negative, positive, uncertainty, litigious, strong model and weak model. In this study, the positive and negative word lists were the only lists that were utilised. The LM lexicon is widely used for text sentiment analysis for financial and economic papers, for example, Cherry and Tong (2023), Shapiro et al. (2019), Shapiro and Wilson (2021), Segawa (2021), and Rutkowska and Szyszko (2024). In the New Zealand context, Cherry and Tong (2023) successfully performed sentiment analysis on monetary policy statements using the LM lexicon (they also utilised the HL lexicon and the Harvard General Inquirer (GI) lexicon). Since the LM lexicon was developed specifically for the fields of economics and finance and has been used in multiple studies, it is the strongest bag-of-words lexicon; therefore, it has been applied in this study. It is an appropriate methodology for comparing the results from the VADER methodology. Refer to Table 2 for the LM descriptive statistics.

Table 2: Loughran McDonald (LM) Descriptive Statistics

Descriptive Statistics	**N**	**Mean**	**Median**	**SD**	**Min**	**Max**
LM Sentiment Score	101	-3.129	-3.000	3.867	-16.000	4.000
LM Sentiment Ratio	101	-0.016	-0.015	0.020	-0.072	0.020
LM Subjectivity Ratio	101	4.983	4.545	1.720	1.770	9.615
LM Polarity Ratio	101	-0.282	-0.333	0.390	-1.000	1.000
Word Count						
LM Positive Word Count	101	3	3	2	0	9
LM Negative Word Count	101	6	6	4	0	16
LM Neutral Word Count	101	177	180	33	94	242
LM Total Word Count	101	186	189	35	98	257

Source: Author's calculations

Table 2 includes descriptive statistics on the following LM variables: LM Sentiment Score, LM Sentiment Ratio, and LM Subjectivity Ratio, as well as the following word count statistics: LM Positive Word Count, LM Negative Word Count, LM Neutral Word Count and LM Total Word Count.

Based on the LM method, in contrast to the VADER sentiment scores, the monetary policy (MP) media releases feature more negative than positive words. On average, there are 3 positive words but 6 negative words per MP media release, resulting in a negative average sentiment score of -3 and a negative sentiment ratio (the sentiment score divided by the average total number of LM tokenised words in the MP media release) of -1.6%. Additionally, the LM subjectivity ratio, which indicates how many positive and negative words there are in a document as a percentage of total LM tokenised words, is ~5%. Put differently, 1 word in 20 is either a positive or negative. Finally, the polarity score, which measures the sentiment score as a percentage of the number of sentiment words, indicates that the average MP media release is neutral, with a polarity score of ~-0.3, which lies within the neutral range between -0.5 and 0.5.

It is important to note that eleven MP media releases have an LM sentiment score of zero (equal number of positive and negative words); therefore, these have LM sentiment ratios and LM polarity ratios of zero (because the numerator in the ratio is zero).

Overall, when the VADER results are compared to the LM results, this study finds, on average, that the RBNZ MP media releases are neutral with a VADER neutral score of 82% and the LM polarity score of –0.3, which lies in the neutral range of –0.5 and 0.5. The LM subjectivity ratio of 5% signals that only 5% of the text carries sentiment, with the other 95% presumably neutral, which is higher than the 82% neutral score of VADER. The LM sentiment ratio demonstrates net negativity as a proportion of total tokenised words. In contrast, VADER had a higher positive score than a negative score, indicating that the positive proportion of the text was greater than the negative. This indicates a slightly conflicting result between the VADER and the LM bag-of-words method. This difference in result could simply be because of the difference in methodologies between VADER, which takes into consideration the whole context of a document.

In contrast, a bag-of-words methodology using a lexicon such as LM is looking at all the words in isolation. Additionally, the LM lexicon was developed to include terminology specifically for finance and economics, whereas VADER was developed to determine sentiment and emotion from social media posts, which were not necessarily specific to economics and finance. However, as per the literature review, VADER has been used successfully to predict the stock market in various studies, for example, Cristescu et al. (2022), Das et al. (2022), Deveikyte et al. (2022), Greyling et al. (2019), Li et al. (2020), and, Soni and Mathur (2023).

In the regression analysis, this study considers the LM sentiment ratio, the LM subjectivity ratio and the LM polarity ratios. All these ratios and the LM lexicon were utilised by the sentiment study on the RBNZ monetary policy statements (MPS) by Cherry and Tong (2023), who determined the sentiment of MPS. This study seeks to investigate whether the LM sentiment of the Monetary Policy (MP) media releases can be used to determine whether an association exists between the sentiment and the stock market and the exchange rate.

3.1.1.2.2 Hu and Liu (HL) covariates

Finally, three sentiment ratios were calculated using the HL lexicon, which has been used in a number of studies and is quite often used in conjunction with the LM lexicon and the GI lexicon (Cherry and Tong, 2023). Hu and Liu (2004) developed their lexicon from customer reviews of specific product features that are listed on e-commerce sites. This process followed three steps: first, they mined for product features that had been commented on by customers; second, they identified opinion sentences in each review and classified them as positive or negative; and third, they summarised the results. Since this lexicon is attached to a tangible positive or negative customer review, this means this lexicon is robust. While the HL lexicon is not based on economic or financial terms like the LM lexicon, it was decided that for robustness purposes, both bag-of-words lexicons will be included to compare the sentiment results. Other studies, such as Shapiro et al. (2020) and Cherry and Tong (2023), included both lexicons to provide more accuracy and expand the audience from people with financial and economic backgrounds to regular people who consume everyday products. Refer to Table 3 for the HL descriptive statistics.

Table 3: Hu and Liu (HL) Descriptive Statistics

Descriptive Statistics	N	Mean	Median	SD	Min	Max
HL Sentiment Score	101	1.475	2.000	4.310	-11.000	10.000
HL Sentiment Ratio	101	0.008	0.010	0.025	-0.069	0.062
HL Subjectivity Ratio	101	8.968	8.805	2.020	2.841	13.636
HL Polarity Ratio	101	0.100	0.111	0.278	-0.846	0.714
Word Count						
HL Positive Word Count	101	9	9	2.973	1	15
HL Negative Word Count	101	7	6	3.441	2	17
HL Neutral Word Count	101	164	168	32.164	83	228
HL Total Word Count	101	180	181	34.982	91	250

Source: Author's calculations

Table 3 includes descriptive statistics on the following HL variables: HL sentiment score, HL sentiment ratio, and HL subjectivity ratio, as well as the following word count statistics: HL positive, HL negative, HL neutral and the HL total word count.

In contrast to the negative average LM sentiment score of -3, the HL sentiment score is net positive, with a mean of 1.5 and a median of 2. The average positive word count is 9, and the average negative word count is 7. This difference could be due to the LM lexicon attributing a negative sentiment to specific words when applied in a financial and economic context, whereas the HL lexicon was created using video product reviews. Therefore, the HL lexicon is more of a generalised lexicon and may not be picking up certain words that carry negativity in an economic context. The HL sentiment ratio (the sentiment score divided by the total number of HL tokenized words) is positive, with an average of ~1% in the MP media releases. Additionally, the HL subjectivity ratio, which indicates the percentage of HL tokenised words carrying sentiment (positive and negative), is 9% on average, which is almost double compared to the LM subjectivity ratio of 5%. Finally, the HL polarity ratio has a mean of 0.10, which falls within the neutrality band of -0.5 to 0.5 on the -1 to 1 scale; this demonstrates text sentiment neutrality for the MP media releases according to the polarity scale.

The HL covariates include the same ratios as the LM lexicon (but with the HL lexicon): the HL sentiment ratio, the HL subjectivity ratio, and the HL polarity ratios.

3.2 Financial and Economic Data

This next section discusses the three financial outcome variables. Our model is inspired by Gorodnichenko et al. (2023), who created a regression model using four different outcome variables (stock market, bond market, inflation expectations and the exchange rate). This study chose to adopt two of these for the New Zealand context. Some of the independent variables Gorodnichenko et al. (2023) included were the sentiment variable, press conference dummy variable and a federal funds rate shock variable. In this study, sentiment covariates are also present, as well as a press conference dummy variable for when there is a televised press conference, and an OCR shock variable has been constructed using the OCR and the overnight index swap (OIS). Refer to the below for more details about each of the variables.

This study investigates the impact of the sentiment within the monetary policy media releases on two markets that have been shown to be responsive to sentiment and are quick to respond to new information: the stock market and the foreign exchange market. Information on these markets, along with covariates, was obtained from Refinitiv Workspace, Bloomberg and the Reserve Bank of New Zealand website.

3.2.1 Dependent Variable 1: New Zealand Stock Exchange (NZX)

For the New Zealand stock exchange data, the daily NZ50E was extracted from Workspace. After extraction, the daily (close to close) returns were calculated; refer to the methodology for further details on the return methodology.

The NZ50E is the S&P/NZX50 Index, which represents the price performance of the 50 largest and most liquid stocks listed on the New Zealand Stock Exchange (NZX). This index (unlike total return indices) does not account for dividends and other distributions. I decided that the NZ50E was the best choice to represent the New Zealand Stock Exchange as it accounts for 90% of New Zealand's market capitalisation (*NZX, New Zealand's Exchange- S&P/NZX Indices*, n.d.). The stock market is open Monday to Friday, 5 days a week; therefore, this study used weekday data. The exchange is closed on public holidays, and this was accounted for in the dataset. The NZ50E will be referred to as the NZX throughout this paper. Refer to Table 4 for the NZX descriptive statistics.

Table 4: NZX Descriptive Statistics

Descriptive Statistics	N	Mean	Median	SD	Min	Max
NZX Day 0 Returns	101	0.0015	0.00106	0.006	-0.0140	0.026
NZX Day 1 Returns	101	0.0008	0.00145	0.006	-0.0131	0.019
NZX Day 2 Returns	101	-0.0001	-0.0001	0.006	-0.0152	0.017
NZX Day 0 to 2 Cumulative Returns	101	0.00214	0.00341	0.011	-0.0257	0.027

Source: Refinitiv Workspace

Based on a 0-to-2-day event window in this study, the following daily returns of the NZX are present: day 0 (the day of the Monetary Policy media release), day 1 and day 2. The mean NZX day 0 returns are 0.15%, day 1 returns are 0.08%, and day 2 returns are -0.01%; therefore, the day 0 to 2 cumulative average returns are 0.214%. This suggests that in the 0 to 2-day event window, the NZX returns are positive on average but with a standard deviation of 0.011, meaning that returns vary. The NZX day 0 to 2 cumulative returns range from -2.6% to 2.7%, therefore a range of over 5%.

3.2.2 Dependent Variable 2: New Zealand Stock Exchange adjusted for offshore markets (NZX SPY)

To account for concurrent events that might have impacted the NZX that are unrelated to the Monetary Policy (MP) media release, the SPDR S&P500 ETF trust data is collected, which is an exchange-traded fund (ETF) tracking the performance of the S&P500 Index in the United States of America (USA). What happens overnight in the USA stock market can impact how the NZ stock market performs the next day. The S&P500 includes the 500 largest and most liquid companies in the USA. The SPDR S&P500 ETF trust will be referred to as the "SPY". Subtracting the SPY return from the NZX gives us a measure of the stock market return adjusted for any major global events that might have impacted stock returns on announcement day. Refer to Table 5 for the NZX SPY descriptive statistics.

Table 5: NZX SPY Descriptive Statistics

Descriptive Statistics	N	Mean	Median	SD	Min	Max
NZX SPY Day 0 Returns	101	0.0007	-0.0007	0.0144	-0.0234	0.1014
NZX SPY Day 1 Returns	101	0.0008	0.0010	0.0114	-0.0487	0.0423
NZX SPY Day 2 Returns	101	-0.0006	-0.0001	0.0125	-0.0576	0.0297
NZX SPY Day 0 to 2 Cumulative Returns	101	0.0010	0.0002	0.0164	-0.0385	0.0500

Source: Refinitiv Workspace

The NZX SPY descriptive statistics are different to the sole NZX returns, as the NZX SPY subtracts the SPY returns from overnight before the NZX opens. In the same format as the NZX Table 4, we have the day 0, 1 and 2 returns individually, as well as the cumulative returns over that same period. The mean NZX SPY day 0 returns are 0.07% (half of the NZX day 0 returns of 0.15%), day 1 returns are 0.08% (the same as the NZX), and day 2 returns are -0.06% (which is larger drop than the NZX day 2 mean returns of -0.001%), therefore the day 0 to 2 cumulative average returns are 0.10%, which is half of the NZX 0 to 2-day cumulative returns of 0.21%. This tells us that in the 0 to 2-day event window, the NZX SPY returns are positive on average but with a standard deviation of 0.016, meaning that returns vary. The NZX SPY day 0 to 2 cumulative returns range from -3.85% to 5%, therefore a range of over 8.85% (1.5 times larger than the NZX alone).

3.2.3 Dependent Variable 3: New Zealand Exchange Rate (TWI)

For the impact on the New Zealand exchange rate data, the daily trade-weighted index (TWI) is extracted from the RBNZ website (RBNZ, n.d). Treasury NZ (2004) defines the TWI as the value of the New Zealand dollar (NZD) against the currencies of 17 currencies representing New Zealand’s major trading partners and is more balanced and less volatile compared to the NZD/USD. The NZD/USD exchange rate was considered, but considering the USA isn’t NZ's largest trading partner and that NZ has many large trading partners, the TWI was a better indicator of the overall effect on the NZ exchange rate. Gorodnichenko et al. (2021), a study based on the federal reserve, used the USD/EUR and the USD/Yen as exchange rate dependent variables. Dowey (2024), who analysed monetary policy communication at a Swedish bank, also used different exchange rates, such as SEK/EUR and SEK/USD. Both studies had conflicting results with each exchange rate that they used, which was probably dependent on how strong they were as a trading partner, along with multiple other factors. For NZ, the TWI is the best way to capture the NZ exchange rates as it is a weighted average of our seventeen largest trading partners. Refer to Table 6 for the TWI descriptive statistics.

Table 6: TWI Descriptive Statistics

Descriptive Statistics	N	Mean	Median	SD	Min	Max
TWI Day 0 Returns	101	-0.0011	0.0000	0.009	-0.0260	0.0180
TWI Day 1 Returns	101	0.0011	0.0007	0.006	-0.0146	0.0133
TWI Day 2 Returns	101	0.0010	0.0004	0.005	-0.0207	0.0145
TWI Day 0 to 2 Cumulative Returns	101	0.0009	0.0020	0.013	-0.0371	0.0311

Source: RBNZ

On the day of the MP media release (day 0), the TWI had a mean return of -0.11%, which is different from the stock market variables, which had positive day 0 returns. The TWI has a mean return of 0.11% for day 1 and a mean return of 0.10% for day 2, and therefore a cumulative 0 to 2 day window mean return of 0.09%. The TWI has a standard deviation of 0.013 for the cumulative 0 to 2-day return window, which is similar and sits in between the NZX (0.0106) and NZX SPY (0.0164). The TWI has a minimum cumulative 0 to 2-day window return of -3.71% and a maximum of 3.11%, therefore showing a range of almost 7%, which is slightly higher than the NZX range of 5% and lower than the NZX SPY range of 8.85%.

Two covariates are used as controls in this study: the press conference dummy variable and the official cash rate surprise variable.

3.2.4 Press conference dummy variable

The press conference dummy variable controls for when the monetary policy (MP) media release is accompanied by a televised press conference (and a full Monetary Policy Statement (MPS)). This binary variable is equal to 1 when there is a press conference and 0 when there is no press conference. This is controlling for the possible impact a speaker (the governor, for example) and a question-and-answer session could have on the stock market or exchange rate. There are four televised press conferences a year (once a quarter), accompanied by a full monetary policy statement (MPS). In one year, there are three MP media releases without a televised press conference; therefore, in total, there are seven MP media releases a year. A basis for controlling for the press conference was based on the study by Gorodnichenko et al. (2021).

3.2.5 OIS OCR - Official cash rate shock variable

An unexpected Official Cash Rate (OCR) change can alter financial markets; therefore, this study needs to control for this. The OCR surprise variable is measured as the OCR on the day of the Monetary Policy (MP) media release, less the close price of the New Zealand Overnight Index Swap (OIS) from the night before. This shock variable will be referred to as the OIS OCR. OIS are interest rate swaps in which the periodic floating rate payment is based on the realised OCR; that is, the observed OIS rate should equal the currently expected OCR over the life of the contract, allowing for market expectations to be observed (Grant & Poskitt, 2024). OIS is the best way to measure the market expectations of the OCR prior to the announcement, and it allows us to account for the effect of unexpected rate changes that are unlikely to have been priced prior to the announcement. The inclusion of the OCR shock is consistent with Gorodnichenko et al. (2023), who also used an interest rate shock variable to control for the impact of an unpredictable change in the federal funds rate. Grant and Poskitt (2024), within the NZ context, also used the OIS to calculate the OCR shock variable. The OIS was retrieved from Bloomberg. Refer to Table 7 for the OIS OCR descriptive statistics.

Table 7: OIS OCR Descriptive Statistics

Descriptive Statistics	N	Mean	Median	SD	Min	Max
OCR	101	3.606	2.500	2.184	1.750	8.250
OCR Change	100	-0.048	0.000	0.253	-1.500	0.250
OIS Daily Rate	101	3.613	2.505	2.193	1.75	8.283
OIS OCR Shock	101	-0.009	-0.003	0.054	-0.295	0.125

Source: RBNZ and Bloomberg

For the 101 observations, the mean OCR is 3.61% (and a median of 2.5%) and ranges from 1.75% to 8.25% (excluding COVID-19, where the OCR neared the lower bound at

0.25%). The mean change of the OCR is close to 0 and ranges from a minimum of decreasing by -1.5% to a maximum of increasing by 0.25%. The OIS has a mean daily rate of 3.61, which is almost identical to the OCR when rounded, and a range between 1.75% and 8.28% (compared to 1% and 8.25%). This OIS OCR shock variable has a mean of -0.009; therefore, on average, it is closely tied to the OCR. The range of shocks is quite large. The OIS OCR shock variable has a minimum of -0.295, demonstrating a situation where the market predicted the OCR to be 29.5 basis points higher than its actual result. This minimum shock occurred on 29 January 2009, when the OCR decreased by 150 basis points, a bigger drop than what the market anticipated; there was no press conference or MPS on this date. The maximum of 0.125 occurred on 7 June 2007, and this was a scenario where the OCR was larger than what the market predicted. This OIS OCR shock maximum occurred on a date when the OCR increased by 25 basis points, and this MP media release was accompanied by an MPS and televised press conference.

Chapter 4 Methodology

This methodology section first discusses the regression analysis to be conducted, which includes an event study to determine whether the sentiment of the monetary policy media releases has an association with the New Zealand stock market and exchange rate. Second, it discusses the OLS diagnostic tests employed to ensure that the OLS conditions hold in the regressions.

4.1 Regression Analysis

To answer the question of whether the sentiment of New Zealand monetary policy media releases is associated with changes in financial indicators (such as the stock market and exchange rate), this study employs an event study and Ordinary Least Square (OLS) regressions. This approach has been applied in previous studies exploring the impact of the sentiment of central bank communications. Shapiro and Wilson (2021) used OLS regressions to determine the central bank's objective function based on sentiment. Gorodnichenko et al. (2021) used panel regressions to explore the impact of sentiment on share prices, while Dowey (2024) used an event study to explore the impact of Riksbank Monetary Policy reports in Sweden on bond market yields and exchange rates.

Hayes (2024) describes an event study as an empirical analysis that examines the impact of a significant catalyst occurrence or contingent event on the value of a security. Event study methodology relies on the foundation of the Efficient Market Hypothesis (EMH). According to the EMH, news affecting economic fundamentals, such as interest rate changes or macroeconomic data releases, should be rapidly incorporated into market prices. Event studies are frequently used in finance to assess the impact of a specific event on a particular stock or market outcome. In this study, the event, the OCR announcement in the form of a Monetary Policy (MP) media release, is examined in relation to how the stock market and exchange rate react to the sentiment of this release. A particular strength of the event study methodology is that it allows for examination of the impact of an event where there may be either speculation or information leakage or where information takes time to be fully or accurately incorporated into markets, such as may be the case with sentiment. Barberis et al. (1997), for instance, argue that sentiment may result in either over or underreaction of stock prices to new information, such as MP media releases. For this study, it was decided that an event study was the best methodology to adopt, considering the subject matter.

In this study, a three-day event window that runs from $t_0 \rightarrow t_{+2}$ was selected (see section 5.3). A shorter window was chosen because longer windows are likely to introduce more variables that could impact the changes in the stock market and exchange rate. For instance, factors such as the release of data from Stats NZ, political announcements such as fiscal policy, or other geopolitical events could introduce noise to the analysis. Dowey (2024) also adopted a shorter event window over three days, and Parle (2022) used a one-day event window. Both these studies referred to the efficient market hypothesis when opting for a shorter event window. Changes prior to the event likely represent speculation rather than information leakage, given the secrecy around OCR announcements (Monetary Policy media releases). In New Zealand, the RBNZ (n.d) requires the Monetary Policy Committee (MPC) to adhere to the MPC charter, which outlines strict rules for outside communication, especially within the two-week window prior to the Monetary Policy Media release. The decision is only made a few days before the announcement, and the charter outlines strict guidelines. It is, therefore, highly unlikely that market participants would be able to get a copy of the release ahead of the announcement, meaning it would not be possible for them to trade on sentiment contained ahead of day 0. Therefore, the three-day event window starting on the day of the announcement (day 0) was chosen.

OLS models are the most common for event studies for several reasons. OLS regressions are straightforward to implement and interpret, making it a clear framework for estimating relationships between variables (Neuhierl et al., 2010).

To determine the returns for the event window, the cumulative returns of day 0, day 1 and day 2 were taken:

$$CR_{t_{0\to 2}} = R_{t_0} + R_{t_1} + R_{t_2} \tag{4}$$

Where $CR_{t_{0\to 2}}$ is the cumulative returns of days 0 to 2, R_{t_0} is the returns on day 0 (the day of the announcement), $R_{t_1}$is the returns one day after the announcement, and R_{t_2} is the returns two days after the announcement. The daily returns are calculated as the log difference between the daily close prices; for example, the returns for day 0 are the log close of day zero minus the log close of day -1.

The cumulative returns for the NZX ($CRNZX_{0\to 2}$) over the 0-to-2-day event window is calculated as follows:

$$\mathrm{CRNZX}_{\mathrm{t}0\to 2} = \sum_{t=0}^{2} \ln(NZXClose_t) - \ln\,(NZXClose_{t-1}) \tag{5}$$

The second dependent variable is the NZX SPY, where the calculation is taken from the NZX above (the log returns from close to close over the three days), and then the log close-to-close SPY returns are subtracted from the NZX. It is common in event studies to account for other events that may occur at the time by computing adjusted returns. In this study, the effect of global events (S&P500, for example) and their impact on the New Zealand stock market (NZX) is considered; therefore, the same-day returns of the SPY are subtracted where:

$$\mathrm{CRSPY}_{\mathrm{t}0\to 2} = \sum_{t=0}^{2} \ln(SPYClose_t) - \ln\,(TWSPYClose_{t-1}) \tag{6}$$

Once we have obtained the cumulative returns for the SPY, we can subtract them from the NZX for the abnormal returns of the NZX SPY: $ARNZXSPY_{t_{0\to 2}}$, where:

$$ARNZXSPY_t = CRNZX_t - CRSPY_t; for\ t_0 \to t_2 \tag{7}$$

Finally, the TWI dependent variable cumulative returns $CRTWI_{t_{0\to 2}}$ is calculated as follows:

$$CRTWI_{t_{0\to 2}} = \sum_{t=0}^{2} \ln(TWIClose_t) - \ln\,(TWIClose_{t-1}) \tag{8}$$

where the $CRTWI$ is the cumulative returns from $\mathrm{t}_0 \to \mathrm{t}_2$, the daily returns of the TWI are the log close minus the log open.

From here on, the NZX will be referred to as the CRNZX, the NZX SPY will be referred to as the ARNZXSPY, and the TWI will be referred to as CRTWI, to be consistent with the methodology.

This study has 24 regressions, eight different types of sentiment regressions: VADER Compound, VADER Positive + VADER Negative, LM Sentiment Ratio, LM Subjectivity Ratio, LM Polarity Ratio, HL Sentiment Ratio, HL Subjectivity Ratio and HL Polarity Ratio for each of the three financial dependent variables: CRNZX, ARNZXSPY and CRTWI.

The generic formula for the six VADER regressions is as follows:

$$Y_t = \beta_0 + \beta_1 X_t + \beta_2 1\{PressConference\}_t + \beta_3 OISOCRShock_t + \epsilon_t \tag{9}$$

Where:

- Y_t: The dependent variable at time t, representing financial indicators such as cumulative returns or abnormal returns (CRNZX, ARNZXSPY, CRTWI).
- β_0: The intercept, indicating the baseline level of Y_t when all independent variables are zero.
- β_1: The coefficient of the sentiment variable (X_t), measuring its effect on Y_t.
- X_t: The VADER sentiment-based variable which can be $VADERCompound\ or\ VADERPositive + VADERNegative$
- β_2: The coefficient for the press conference dummy indicator variable, measuring the effect of monetary policy press conferences.
- β_3: The coefficient for the OISOCRShock, capturing the effect of surprises in the Official Cash Rate (OCR).

The generic formula for the nine Loughran McDonald (LM) regressions is as follows:

$$Y_t = \beta_0 + \beta_1 X_t + \beta_2 1\{PressConference\}_t + \beta_3 OISOCRShock_t + \epsilon_t \tag{10}$$

Where:

- Y_t: The dependent variable at time t, representing financial indicators such as cumulative returns or abnormal returns (CRNZX, ARNZXSPY, CRTWI).

- β_0: The intercept, indicating the baseline level of Y_t when all independent variables are zero.
- β_1: The coefficient of the sentiment variable (X_t), measuring its effect on Y_t.
- X_t: The LM sentiment-based variable which can be $LMSentimentRatio, LMSubjectivityRatio\ or\ LMPolarityRatio$
- β_2: The coefficient for the press conference dummy indicator variable, measuring the effect of monetary policy press conferences.
- β_3: The coefficient for the OISOCRShock, capturing the effect of surprises in the Official Cash Rate (OCR).

The generic formula for the nine Hu and Liu (HL) regressions is as follows:

$$Y_t = \beta_0 + \beta_1 X_t + \beta_2 1\{PressConference\}_t + \beta_3 OISOCRShock_t + \epsilon_t \quad (11)$$

Where:

- Y_t: The dependent variable at time t, representing financial indicators such as cumulative returns or abnormal returns (CRNZX, ARNZXSPY, CRTWI).
- β_0: The intercept, indicating the baseline level of Y_t when all independent variables are zero.
- β_1: The coefficient of the sentiment variable (X_t), measuring its effect on Y_t.
- X_t: The HL sentiment-based variable which can be $HLSentimentRatio, HLSubjectivityRatio\ or\ HLPolarityRatio$
- β_2: The coefficient for the press conference dummy indicator variable, measuring the effect of monetary policy press conferences.
- β_3: The coefficient for the OISOCRShock, capturing the effect of surprises in the Official Cash Rate (OCR).

To ensure the robustness of the results, OLS diagnostic tests were conducted. These are outlined in the next section.

4.2 Diagnostic Tests

This section reports on the diagnostic tests conducted to ensure that the OLS conditions hold for the regressions in this study. Specifically, the six key OLS assumptions outlined by Wooldridge (2019) are assessed: linearity in parameters, no perfect collinearity, zero conditional mean, homoscedasticity, no serial correlation, and normality and independence of residuals. These tests aim to identify potential issues requiring correction. A detailed explanation of each assumption and the corresponding diagnostic tests follows.

4.2.1 Linearity in Parameters

The first assumption requires that the relationship between the dependent variables (returns of the CRNZX, ARNZXSPY, and CRTWI) and the independent variables (9 sentiment variables, press conference dummy variable, and OCR shock variable) is linear in parameters.

To test for linearity, this study employed the Ramsey RESET test for omitted variables. The hypotheses are:

- Null hypothesis (H_0): The model is correctly specified and linear in parameters.
- Alternative hypothesis (H_1): The model is misspecified, potentially due to nonlinearity or omitted variables.

The *p*-values are interpreted as follows:

- *p*<0.1: Reject H_0, indicating misspecification.
- *p*>0.1: Fail to reject H_0, indicating the model is correctly specified.

Nineteen of the twenty-four regressions satisfied the RESET test. None of the VADER specifications failed, while the CRNZX failed when regressed against sentiment scores and subjectivity ratios using either the LM or HL dictionaries. Refer to Appendices A, B, and C for detailed results.

4.2.2 No Perfect Collinearity

The second assumption requires that no independent variable is constant or a perfect linear combination of others (Wooldridge, 2019). Perfect collinearity prevents the unique estimation of coefficients and renders OLS computation infeasible.

To test for multicollinearity, the Variance Inflation Factor (VIF) was calculated using the estat vif command in Stata. A VIF score:

- <10: Multicollinearity is not a concern.
- ≥10: Suggests strong collinearity and violates the assumption.

All 24 regressions produced VIF scores between 1 and 1.07, indicating no significant multicollinearity. Results are available in Appendices A, B and C.

4.2.3 Zero Conditional Mean (Exogeneity)

The zero conditional mean assumption ensures that independent variables are uncorrelated with the error term across all time periods. Specifically, the error term at time $t(\epsilon_t)$ must be uncorrelated with each explanatory variable in all periods (Wooldridge, 2019).

This study assessed exogeneity through two methods:

1. Ramsey RESET test for omitted variable bias, as described above.
2. Calculation of correlations between residuals and independent variables.

No significant correlations were detected between residuals and independent variables across all 24 regressions, affirming the exogeneity assumption. A detailed correlation matrix is provided in Appendix D.

4.2.4 Homoscedasticity

Homoscedasticity assumes constant variance of errors $(Var(\epsilon_t|X_t) = \sigma^2)$. Violations (heteroscedasticity) yield invalid standard errors and hypothesis tests (Wooldridge, 2019).

The following tests were conducted:

1. White's test (estat imtest, white): Null hypothesis (H_0): Homoscedasticity.
2. Cameron and Trivedi's (C&T) Decomposition of IM Test: Provides insights on heteroscedasticity, skewness, kurtosis, and an overall score.

Interpretation of *p*-values:

- *p*<0.1: Reject H_0, indicating heteroscedasticity.
- p>0.1: Fail to reject H_0, suggesting homoscedasticity.

All 24 regressions passed the White test; 23 passed the C&T test. Refer to Appendices A, B, C and D for the full test results.

4.2.5 No Serial Correlation

Serial correlation is a violation of the independence of errors; when this test fails, it indicates that the standard errors are biased, and this affects inferences. No serial correlation implies uncorrelated errors across time $(Cov(\epsilon_t, \epsilon_{t-s}|X_t) = 0, \forall_s \neq 0)$.

This assumption was tested by:

1. Predicting residuals using *predict resid, residuals*.
2. Generating lagged residuals: *gen resid_lag = resid[_n-1]*.
3. Regressing residuals on lagged residuals: *regress resid resid_lag*.

The hypotheses are:

- H_0: No serial correlation.
- H_1: Serial correlation exists.

P-value interpretation:

- p<0.1: Reject H_0, indicating serial correlation.
- p>0.1: Fail to reject H_0, indicating no serial correlation.

If all five aforementioned OLS assumptions hold, this confirms the Gauss-Markov Theorem that the OLS estimators are the best linear unbiased estimators conditional on X (Wooldridge, 2019). All regressions produced high *p*-values, affirming no serial correlation. The results are in Appendices A, B, C and D.

4.2.6 Normality and Independence of Residuals

The final assumption requires errors to be normally and independently distributed $\left(\epsilon_t \sim N(0, \sigma^2)\right)$ (Wooldridge, 2019). The Shapiro-Wilk test (swilk residuals) was adopted to assess normality:

- H_0: Errors are normally distributed.
- H_1: Errors are not normally distributed.

P-value interpretation:

- p<0.10: Reject H_0, indicating non-normality.
- p>0.10: Fail to reject H_0, indicating normality.

Given the Central Limit Theorem, large sample sizes (*n*>100) mitigate concerns about non-normality. All 24 regressions produced *p*>0.10, confirming normality. Refer to Appendices A, B, and C for the full results.

As a result, strong evidence is found to suggest that OLS regressions will give unbiased estimates of the relationship between sentiment and market outcomes.

Chapter 5 Results and Discussion

This section discusses the sentiment analysis results, including the key findings from the descriptive statistics and how the different sentiment lexicons VADER, Loughran McDonald (LM) and Hu and Liu (HL) differ in results. The main part of this section will discuss the regression analysis results across the different dependent variables, the cumulative returns of the New Zealand Stock Exchange Top 50 Index (CRNZX), CRNZX minus the USA's S&P500 Index (CRSPY), which forms the ARNZXSPY, and the trade-weighted index (CRTWI), and the three different sentiment analyses. Refer to Appendix E for regression coefficients and statistical significance. Lastly, I will discuss the robustness checks, limitations and extensions for future research. I begin by examining the regression results of our primary sentiment technique, VADER.

5.1 VADER Regression Analysis Results

Across all three dependent variables, all VADER regressions passed the OLS diagnostic tests. As such, I conduct six VADER regressions, specifically for each of the three dependent variables. Refer to equation (9) in section 4.1.

Across CRNZX, ARNZXSPY, and CRTWI, there were no significant results in the VADER compound score regressions, indicating that the VADER compound score was not significantly related to subsequent movements in the stock market or exchange rate.

In regard to the other VADER regressions, including the VADER positive and negative score, there were no significant sentiment coefficients when regressed against the CRNZX, but there was a marginally significant coefficient in the ARNZXSPY – refer to Table 8. This is to be expected as the ARNZXSPY accounts for global effects that might have impacted the NZ market, effectively removing outside events from the analysis. In Table 8 for the ARNZXSPY, the VADER negative score had a coefficient of -0.140, which was significant at the 10% level. A negative association between the stock market and the VADER negative score indicates that a one-unit increase in the VADER negative score could lead to a 0.14 unit decrease in the log returns (or a ~13% decrease) of the

ARNZXSPY. This suggests that greater negative sentiment within an announcement, on average, results in a decline in the stock market. This is in line with expectations, where negative sentiment about the economy in the future could lead to decreased returns. This same regression had a low R-squared, showing 4.3%, indicating that this very simple specification only explains about 4% of the observed variance in the ARNZXSPY. However, after adjusting for a number of predictors, the adjusted R-squared shows almost no variability in the ARNZXSPY, which is accounted for by this model. The model wasn't statistically significant overall, with a Prob > F of 0.3684.

Table 8: VADER Highlighted Significant Results

Dependent Variable	Sentiment Variable(s)	R Squared	Adjusted R Squared	Significant Coefficient	Prob > F
ARNZXSPY	VADER Positive	0.043	0.0034		0.3684
	VADER Negative	0.043	0.0034	**-0.140***	0.3684
CRTWI	VADER Positive	0.18	0.146	**-0.0917****	0.0007
	VADER Negative	0.18	0.146		0.0007

Source: Author's calculations
Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The final VADER regression, with the positive and negative regressors, has the CRTWI as the dependent variable. In this regression, the VADER positive regressor was statistically significant and negative, showing a negative association between the positive sentiment score and the CRTWI returns at the 5% level of significance. A one-unit increase in the VADER positive score could lead to a 0.0917 decrease in the log returns (or a ~9.6% decrease) of the CRTWI. This is the opposite of the ARNZXSPY result, which showed a significant (p value<0.1) negative association between the VADER negative score and the dependent variable returns. This result makes logical sense. This stronger negative association between the positive score and the CRTWI suggests that optimistic monetary policy releases might signal to markets that the RBNZ is adopting a dovish stance (e.g., signalling low rates or economic challenges ahead), which could reduce capital inflows, depreciating the New Zealand Dollar. This final VADER CRTWI regression has an R-squared that shows that 18% of the variance of the

CRTWI can be explained by this regression model, which is still low but close to four times larger than the R-squared ARNZXSPY regression. Overall, for the CRTWI, the VADER positive and negative regression equation has a Prob > F of 0.0007 and is statistically significant.

5.2 Loughran McDonald (LM) Regression Analysis Results

The LM regressions did not pass the diagnostic tests as well as the VADER regressions. I found that three of the nine regressions failed the RESET omitted variables. Two of these were CRNZX regressions, one with the LM Sentiment regressor and the other with the LM Subjectivity regressor – refer to equation (10) in section 4.1. The other regression which failed was the ARNZXSPY regression with the LM Subjectivity regressor. Refer to Appendix B for LM OLS diagnostic tests. The ARNZXSPY LM Subjectivity regression will not be interpreted because it failed this test.

In addition to the above regressions, which failed the RESET test, I also failed to find significant associations between the CRNZX and ARNZXSPY for any of the LM sentiment measures: Sentiment, Subjectivity and Polarity (except ARNZXSPY and LM Subjectivity, which will not be interpreted as mentioned above).

All three of the CRTWI LM regressions passed the OLS testing, and two of these CRTWI regressions (LM Sentiment and LM Polarity) showed significant relationships between the LM sentiment regressors and the CRTWI.

Table 9: LM Highlighted Significant Results

Dependent Variable	Sentiment Variable(s)	R Squared	Adjusted R Squared	Significant Coefficient	Prob > F
CRTWI	LM Sentiment Ratio	0.164	0.138	0.138**	0.0006
CRTWI	LM Subjectivity Ratio	0.161	0.1346	-0.00148**	0.0007

Source: Author's calculations

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 9 shows that for the first LM CRTWI regression, there is a positive association between the LM sentiment ratio regressor and the trade-weighted index at the 5% level of significance. A one-unit increase in the LM sentiment ratio could lead to a 0.138 unit increase (or a ~14.8% increase) in the CRTWI log returns on average. This is contrary to the findings earlier mentioned in the CRWTI VADER regression and possibly needs further investigation. The LM sentiment ratio is the number of positive words minus the number of negative words divided by the number of LM tokenised words in the Monetary Policy (MP) media release. Therefore, a more net positive LM sentiment ratio has a positive effect on the CRTWI daily returns. On average, this ratio was net negative (more negative words than positive) and contrary to both the VADER and HL measures, which were net positive. According to the R-squared for this regression, approximately 16.4% of the variability in the CRTWI is explained by the LM Sentiment Ratio. It has a Prob > F of 0.0006; therefore, this model is statistically significant. The TWI VADER Positive/Negative regression has a marginally higher R-squared value than the LM sentiment regression, but the difference is small, and both models provide significant insights into CRTWI variability.

Table 9 shows that for the CRTWI LM subjectivity regression, a negative association between the CRTWI and the LM subjectivity score regressor exists at the 5% level of significance. The LM subjectivity ratio is the number of words that carry sentiment (positive and negative) divided by the number of LM tokenised words. A one-unit increase in the LM subjectivity score could lead to a 0.00148 decrease (or a 0.15% decrease) in the CRTWI; although statistically significant, economically, this is small in magnitude. This shows us that the higher LM Subjectivity Score (reflecting greater subjectivity in sentiment) is significantly but weakly associated with a decrease in the CRTWI. However, it is worth noting that the increase in subjectivity could be due to either an increase in positive or negative words or a combination of both. The R-square shows that approximately 16.1% of the variability in the CRTWI is explained by the LM Subjectivity Ratio, which is slightly weaker than the 16.4% in the former LM CRTWI regression. However, with a Prob > F of 0.0007, this CRTWI LM subjectivity regression model is also statistically significant overall.

Overall, LM sentiment appears to be associated with changes in the cumulative log returns of the trade-weighted index (CRTWI) as all the OLS tests passed, and there were two regressions with significant sentiment regressors. These models didn't work well for the stock market dependent variables CRNZX and ARNZXSPY as there was omitted variable bias in three of the regressions, and the stock market regressions that passed all the tests didn't have a significant relationship between the stock market and the LM sentiment regressors. VADER seems to be slightly better at measuring the relationship between sentiment in MP media releases and the stock market (shown in the ARNZXSPY VADER positive negative regression). For the exchange rate, both the LM and VADER models show valuable insights into the association between sentiment and the exchange rate (CRTWI VADER positive negative and CRTWI LM Sentiment regression).

5.3 Hu and Liu (HL) Regression Analysis Results

The final group of regressions are the HL models. I did not anticipate finding significant results in the HL regressions because the lexicon was not finance- or economic-specific. However, this study found one regression with significant results; refer to Table 10.

Table 10: HL Highlighted Significant Results

Dependent Variable	Sentiment Variable(s)	R Squared	Adjusted R Squared	Significant Coefficient	Prob > F
TWI	HL Subjectivity Ratio	0.166	0.1397	-0.00135**	0.0006

Source: Author's calculations
Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Again, some of the HL regressions failed the OLS tests, similar to the three in the LM regressions that failed the RESET omitted variables test. Two CRNZX HL regressions (HL Sentiment and HL Polarity) failed the RESET omitted variables test, while the

CRNZX HL Sentiment regression also failed the Cameron and Trivedi IM breakdown. The other seven HL regressions passed all the OLS diagnostic tests.

The only HL regression with significant results was the CRTWI HL Subjectivity regression – refer to equation (11) in section 4.1. This study found a negative association between the HL subjectivity score regressor and the trade-weighted index at the 5% level of significance (Table 10). A one-unit increase in the HL subjectivity score could result in a 0.00135 unit decrease in the log returns (or a 0.14% increase) of CRTWI. This is Similar to the LM subjectivity regressor, which also had an economically small negative association with the CRTWI. The HL subjectivity CRTWI regression has a slightly higher R-squared than the LM subjectivity regression (0.164 vs 0.166). This CRTWI HL sentiment regression shows that 16.6% of the variation in the CRTWI can be explained by this model, and with a Prob > F of 0.0006, this model is statistically significant overall.

5.3 Robustness

To ensure the robustness of the results, multiple regression models were run using different sentiment measures and time windows. Initially, I conducted regression analysis on a -5 to 5-day window, a 4- to 4-day window, a -3 to 3-day window, a -2 to 2-day window, a -1 to 1-day window, 0 to 1-day window, 0 to 2-day window (our current model), 0 to 3-day window, 0 to 4-day window and a 0 to 5-day window. While I believed that the 0-2 day window was the most appropriate, I also tested a range of additional event windows to ensure that this was the case. Specifically, I tested a range of windows beginning before the announcement day to control for potential information leakages that might have meant not all of the reaction was accounted for in the 0-2 day window. Additionally, the longer event windows were used in case the New Zealand financial markets were less efficient than assumed and the market continued to adjust for changes during the following business week. I found that event windows that began on day 0 generated the strongest results, with most OLS tests passed, but as the window increased past the 0-to-2-day mark, more tests started failing. As noted above in section 4.1, there is unlikely to be information leakage prior to the Monetary Policy (MP) media releases, meaning any pre-event movements are likely to be driven by speculation or other events. Additionally, 0 to 2 days post-event covers trading to the end of the business week, with longer event windows beginning to incorporate more unrelated information.

Three sentiment methodologies were applied to determine the sentiment of the MP media releases to ensure the robustness of the sentiment measure. The main focus was on the comparison between VADER and the LM bag-of-words lexicon, with the hypothesis that VADER would generate more significant results. The HL bag-of-words lexicon was used as a general lexicon to add robustness to the LM lexicon. By utilising three lexicons, I can be confident that the sentiment analysis applied was accurate and robust.

To determine whether the sentiment of monetary policy media releases impacted financial indicators, I utilised three dependent variables to ensure robustness for the stock market: CRNZX and ARNZXSPY and CRTWI for the exchange rate. The ARNZXSPY variable was a more robust indicator of the New Zealand stock market; it generated interpretable results given that it passed all of the OLS tests compared to the CRNZX alone, this being due to the fact it accounts for the performance of the US stock market overnight. By applying three dependent variables to separate models, I can draw conclusions about the stock market and exchange rate and analyse New Zealand's financial markets in general.

I made sure that the necessary OLS diagnostic tests were applied to ensure the regressions held up. I also ran regressions with robust standard errors; this helped some of the results, but it did not help improve the models that had omitted variable bias.

Chapter 6 Conclusion

This study aimed to examine whether the text sentiment of New Zealand Monetary Policy (MP) media releases impacts financial indicators, specifically the stock market and exchange rate. To do this, I took Cherry and Tong's (2023) study on the sentiment of monetary policy statements and also applied both the LM and HL methodologies as well as VADER. Additionally, I applied this to MP media releases rather than the monetary policy statements, which only occur 4 times per year. Greyling et al. (2019) had success when using VADER to predict the stock market, and Shapiro and Wilson (2021) used VADER to perform sentiment analysis on Federal Reserve monetary policy documentation to predict a central bank loss function successfully. I hypothesised that VADER could generate better, if not valuable, results compared to the bag-of-words methods because it analyses the sentiment of complete sentences rather than the count of words in the monetary policy media releases that appear in a lexicon.

My sentiment results show that all three sentiment methodologies determined the MP media releases were mostly neutral on average, which is consistent with Cherry and Tong’s (2023) findings. However, VADER appeared to find slightly more positive than negative sentiment within the MP media releases. The LM method had a net negative sentiment score and ratios, while the HL was also positive but very close to zero with its sentiment scores and ratios, reiterating the neutrality of the RBNZ monetary policy media releases.

My regression analysis shows that, indeed, VADER produced valuable results where all six regressions passed all the OLS diagnostic tests and provided the only group of regressions with a significant relationship between the exchange rate and the New Zealand stock market. In the ARNZXSPY VADER positive and negative sentiment regression, there was a marginally negative association between the negative VADER score covariate and the ARNZXSPY, showing that a higher proportion of negative sentiment in the MP media releases could negatively impact the New Zealand stock market. In the CRTWI VADER positive and negative regression, there was a more significant negative association between the positive score covariate and the cumulative log returns of the trade-weighted index, where an increase in positive sentiment as a proportion of MP media release could negatively impact the New Zealand exchange rate.

The results of the bag of words regressions were more mixed. Some regressions failed OLS diagnostic tests, and there was no association between any of the sentiment variables and the stock market variables (CRNZX and ARNZXSPY). The CRTWI regressions had the best results, with a significant association between the LM sentiment ratio, LM subjectivity ratio, and HL subjectivity ratio. In this group of regressions, the most important ratio is the LM sentiment ratio, which can be determined as net positive or net negative sentiment. In comparison, the subjectivity ratios capture the proportion of sentiment generally (positive and negative) as a proportion of the entire text.

Overall, across all three sentiment methodologies and regressions, I found a statistically significant association between the sentiment of the MP media releases and the exchange rate, but at best, it was a marginal association with the stock market in the case of the VADER positive negative ARNZXSPY regression. Overall, VADER produced slightly better and more interpretable results than the LM methods due to the VADER regressions passing all of the OLS tests, performing just as well as the LM as having statistically and economically significant results for the exchange rate, as well as having interpretable results for the stock market which the LM methods did not have. My study confirms that the way information is communicated and the sentiment it carries can potentially impact financial markets, specifically in the case of central bank monetary policy communications. This highlights the importance for policymakers in the choice of language when making communicating public, and specifically monetary, policy. Additionally, if used purposefully, policymakers could use the sentiment contained in their announcements to either reaffirm or mitigate reactions to the rate decision.

While this study provides valuable insights, it also has some limitations. Daily data was used for the exchange rate and stock market, which may not fully capture the intraday market reactions. Additionally, this study only focused on the stock market and exchange rate rather than the plethora of measures available in the financial markets. Future research could explore intraday data to capture more granular market responses and include other dependent variables such as treasury bills, bonds and inflation. My model, while based on Gorodnichenko et al.'s (2021) model, does not control for forward guidance, large-scale asset purchases, the shadow rate and voice tone. Voice tone is the focal point in Gorodnichenko et al.'s (2021) study. It would be beneficial to extend our study to control for the introduction of the monetary policy committee and the period during COVID-19 and beyond so that a larger date range can be utilised. Ideally, adopting Swanson's (2021) three-factor model, which was utilised in Gorodnichenko et al.'s (2021) study to control for forward guidance, large-scale asset purchases, which data was not available, and the federal funds rate shock – which I included via the OIS OCR shock variable. It would also be interesting to see the sentiment of the RBNZ's televised press conferences in New Zealand and how the sentiment differs across different governors and different media correspondents; these all form part of a future research agenda.

By uncovering the association between the sentiment of monetary policy media releases and financial indicators, this study contributes to a growing body of research on the intersection of central bank communications and financial markets. It contributes to the development of more precise frameworks centred around policy transparency.

Chapter 7 References

Abercrombie, G., & Batista-Navarro, R. (2020). Sentiment and position-taking analysis of parliamentary debates: a systematic literature review. *Journal of Computational Social Science*, *3*(1), 245–270. https://doi.org/10.1007/s42001-019-00060-w

Akladyous, B. (2023, October 23). Sentiment Analysis using VADER - Boula Akladyous - Medium. *Medium*. https://akladyous.medium.com/sentiment-analysis-using-vader-c56bcffe6f24

Almeida, M. D., Maia, V. M., Tommasetti, R., & De Oliveira Leite, R. (2021). Sentiment analysis based on a social media customised dictionary. *MethodsX*, *8*, 101449. https://doi.org/10.1016/j.mex.2021.101449

Anderson, J. C., & Thorshaug, M. G. (2021). *Central Bank Sentiment Analysis and Asset Prices* [Masters Thesis]. Norwegian School of Economics.

Barberis, N., Shleifer, A., & Vishny, R. (1997). *A model of investor sentiment*. https://doi.org/10.3386/w5926

Cherry, R., & Tong, E. (2023). *Words of the RBNZ*. Reserve Bank of New Zealand. https://www.rbnz.govt.nz/hub/publications/analytical-note/2023/an-2023-03

Cristescu, M. P., Nerisanu, R. A., Mara, D. A., & Oprea, S. (2022). Using market news sentiment analysis for stock market prediction. *Mathematics*, *10*(22), 4255. https://doi.org/10.3390/math10224255

Das, N., Sadhukhan, B., Chatterjee, T., & Chakrabarti, S. (2022). Effect of public sentiment on stock market movement prediction during the COVID-19 outbreak. *Social Network Analysis and Mining*, *12*(1). https://doi.org/10.1007/s13278-022-00919-3

Deveikyte, J., Geman, H., Piccari, C., & Provetti, A. (2022). A sentiment analysis approach to the prediction of market volatility. *Frontiers in Artificial Intelligence*, *5*. https://doi.org/10.3389/frai.2022.836809

Devlin, J., Chang, M., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *Association for Computational Linguistics*, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. https://doi.org/10.18653/v1/n19-1423

Dowey, L. (2024). *Unveiling Economic Sentiment: Using a Large Language Model for Economic Sentiment Analysis from Monetary Policy Reports* [Masters Thesis]. CHALMERS UNIVERSITY OF TECHNOLOGY UNIVERSITY OF GOTHENBURG.

GeeksforGeeks. (2024, January 31). *NLP | How tokenizing text, sentence, words works*. GeeksforGeeks. https://www.geeksforgeeks.org/nlp-how-tokenizing-text-sentence-words-works/

Gorodnichenko, Y., Pham, T., & Talavera, O. (2023). The voice of monetary policy. *American Economic Review*, *113*(2), 548–584. https://doi.org/10.1257/aer.20220129

Grant, R., & Poskitt, K. (2024). *Financial Market Reaction to Monetary Policy Surprises*. Reserve Bank of New Zealand. https://www.rbnz.govt.nz/-/media/project/sites/rbnz/files/publications/analytical-notes/2024/financial-market-reaction-to-ocr-surprises.pdf

Rossouw, S., & Greyling, T. (2020). Big Data and Happiness. In *Springer eBooks* (pp. 1–35). https://doi.org/10.1007/978-3-319-57365-6_183-1

Greyling, T., Rossouw, S., & Steyn, D. H. W. (2022). The Prediction of Intraday Stock Market Movements in Developed & Emerging Markets using Sentiment and Emotions from Twitter. *Indian Institute of Finance*, *XXXVI*(3), 907–939.

Hailong, Z., Wenyan, G., & Bo, J. (2014). Machine Learning and Lexicon based Methods for Sentiment Classification: A Survey. *In 2014 11th Web Information System and Application Conference*. https://doi.org/10.1109/wisa.2014.55

Harvard. (n.d.). *Harvard General Inquirer*. Harvard General Inquirer Home Page. https://inquirer.sites.fas.harvard.edu/

Hayes, A. (2024, August 6). *Event Study: Definition, Methods, Uses in Investing and Economics*. Investopedia. https://www.investopedia.com/terms/e/eventstudy.asp#:~:text=An%20event%20study%20is%20an,security%2C%20such%20as%20company%20stock.

Hu, M., & Liu, B. (2004). *Mining and summarizing customer reviews*. https://doi.org/10.1145/1014052.1014073

Hutto, C., & Gilbert, E. (2014). VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text. *Proceedings of the International AAAI Conference on Web and Social Media*, *8*(1), 216–225. https://doi.org/10.1609/icwsm.v8i1.14550

Li, X., Wu, P., & Wang, W. (2020). Incorporating stock prices and news sentiments for stock market prediction: A case of Hong Kong. *Information Processing & Management*, *57*(5), 102212. https://doi.org/10.1016/j.ipm.2020.102212

Loughran, T., & Mcdonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, *66*(1), 35–65. https://doi.org/10.1111/j.1540-6261.2010.01625.x

Mohammad, S., & Turney, P. (2010). Emotions evoked by common words and phrases: using mechanical Turk to create an emotion lexicon. *North American Chapter of the Association for Computational Linguistics*, 26–34. https://saifmohammad.com/WebDocs/Mohammad-Turney-NAACL10-EmotionWorkshop.pdf

Neuhierl, A., Scherbina, A., & Schlusche, B. (2010). Market reaction to corporate press releases. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.1556532

Nguyen, B., & Huynh, V. (2020). Textual analysis and corporate bankruptcy: A financial dictionary-based sentiment approach. *Journal of the Operational Research Society*, *73*(1), 102–121. https://doi.org/10.1080/01605682.2020.1784049

NZX, New Zealand's Exchange- S&P/NZX indices. (n.d.). https://www.nzx.com/markets/indices

Parle, C. (2022). The financial market impact of ECB monetary policy press conferences — A text based approach. *European Journal of Political Economy*, *74*, 102230. https://doi.org/10.1016/j.ejpoleco.2022.102230

RBNZ. (n.d.). *Exchange rates and Trade Weighted Index (B1)*. Reserve Bank of New Zealand. https://www.rbnz.govt.nz/statistics/series/exchange-and-interest-rates/exchange-rates-and-the-trade-weighted-index

Rutkowska, A., & Szyszko, M. (2024). Dictionary-based sentiment analysis of monetary policy communication: on the applicability of lexicons. *Quality & Quantity*. https://doi.org/10.1007/s11135-024-01896-9

Segawa, A. (2021). Sentimental outlook for the monetary policies of South African Reserve Bank. *International Journal of Finance & Banking Studies (2147-4486)*, *10*(3), 37–56. https://doi.org/10.20525/ijfbs.v10i3.1298

Shahzad, S. J. H., Kyei, C. K., Gupta, R., & Olson, E. (2020). Investor sentiment and dollar-pound exchange rate returns: Evidence from over a century of data using a cross-quantilogram approach. *Finance Research Letters, 38*, 101504. https://doi.org/10.1016/j.frl.2020.101504

Shapiro, A. H., & Wilson, D. J. (2021). Taking the Fed at its Word: A New Approach to Estimating Central Bank Objectives using Text Analysis. *The Review of Economic Studies*, *89*(5), 2768–2805. https://doi.org/10.1093/restud/rdab094

Shapiro, A., Wilson, D., & Sudhof, M. (2019). Measuring News Sentiment - San Francisco Fed. *SF Fed*. https://www.frbsf.org/research-and-insights/publications/working-papers/2020/03/measuring-news-sentiment/

Soni, J., & Mathur, K. (2023). Sentiment Analysis of News Headlines for Stock Market Prediction using VADER. In *International Conference on Innovative Mechanisms for Industry Applications (ICIMIA)* (pp. 1215–1222). https://doi.org/10.1109/icimia60377.2023.10426095

Swanson, E. T. (2020). Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. *Journal of Monetary Economics*, *118*, 32–53. https://doi.org/10.1016/j.jmoneco.2020.09.003

The Treasury NZ. (2004, January 9). *Aide Memoire: Recent trends in the New Zealand dollar exchange rate*. The Treasury New Zealand. https://www.treasury.govt.nz/publications/aide-memoire-recent-trends-new-zealand-dollar-exchange-rate

Turney, & Mohammad. (2010). North American Chapter of the Association for Computational Linguistics. In *Emotions evoked by common words and phrases: using mechanical Turk to create an emotion lexicon* (pp. 26–34). https://saifmohammad.com/WebDocs/Mohammad-Turney-NAACL10-EmotionWorkshop.pdf

Van Dijk, D., Lumsdaine, R. L., & Van Der Wel, M. (2016). Market set-up in advance of federal reserve policy rate decisions. *The Economic Journal, 126*(592), 618–653. https://doi.org/10.1111/ecoj.12372

Wooldridge, J. M. (2019). *Introductory Econometrics: a Modern Approach* (7th ed.). http://ci.nii.ac.jp/ncid/BA72932652

Zeng, X., Sun, Y., Xu, Y., & Hu, H. (2025). Monetary Policy Shocks, Investor Sentiment, and Stock Returns: Firm-Level Evidence from China. *Emerging Markets Finance and Trade*, 1–18. https://doi.org/10.1080/1540496X.2025.2467819

Appendix A: VADER Summary of OLS Diagnostic Results

Dependent Variable	Sentiment Variable(s)	RESET Omitted Variables	VIF Multicollinearity	Uncorrelated Residuals	White Homoskedasticity	Cameron & Trivedi IM	No Serial Correlation	SWILK Normality	Overall
CRNZX	VADER Compound	P 0.2067	P 1.01	P	P 0.1645	P 0.1341	P 0.957	P 0.58425	P
ARNZXSPY	VADER Compound	P 0.1000	P 1.01	P	P 0.8520	P 0.8467	P 0.156	P 0.39696	P
CRTWI	VADER Compound	P 0.8749	P 1.01	P	P 0.3628	P 0.4928	P 0.831	P 0.27199	P
CRNZX	VADER Positive VADER Negative	P 0.2473	P 1.03	P	P 0.8943	P 0.7599	P 0.968	P 0.57112	P
ARNZXSPY	VADER Positive VADER Negative	P 0.2206	P 1.03	P	P 0.3666	P 0.5439	P 0.249	P 0.58308	P
CRTWI	VADER Positive VADER Negative	P 0.5483	P 1.03	P	P 0.2493	P 0.4121	P 0.657	P 0.95490	P

Key: P= PASS, F = FAIL

Appendix B: LM Summary of OLS Diagnostic Results

Dependent Variable	Sentiment Variable(s)	RESET Omitted Variables	VIF Multicollinearity	Uncorrelated Residuals	White Homoskedasticity	Cameron & Trivedi IM	No Serial Correlation	SWILK Normality	Overall
CRNZX	LM Sentiment Ratio	F 0.0252	P 1.07	P	P 0.8229	P 0.5488	P 0.704	P 0.58651	F
ARNZXSPY	LM Sentiment Ratio	P 0.3896	P 1.07	P	P 0.8972	P 0.8058	P 0.200	P 0.25108	P
CRTWI	LM Sentiment Ratio	P 0.7846	P 1.07	P	P 0.3538	P 0.2959	P 0.200	P 0.48918	P
CRNZX	LM Subjectivity Ratio	F 0.0416	P 1.02	P	P 0.8396	P 0.6840	P 0.744	P 0.70783	F
ARNZXSPY	LM Subjectivity Ratio	F 0.0973	P 1.02	P	P 0.7855	P 0.6643	P 0.241	P 0.62366	F
CRTWI	LM Subjectivity Ratio	P 0.6551	P 1.02	P	P 0.3484	P 0.3973	P 0.968	P 0.71739	P
CRNZX	LM Polarity Ratio	P 0.4876	P 1.06	P	P 0.6785	P 0.4544	P 0.733	P 0.56088	P
ARNZXSPY	LM Polarity Ratio	P 0.5340	P 1.06	P	P 0.7741	P 0.7727	P 1.167	P 0.31022	P

Appendix C: HL Summary of OLS Diagnostic Results

Dependent Variable	Sentiment Variable(s)	RESET Omitted Variables	VIF Multicollinearity	Uncorrelated Residuals	White Homoskedasticity	Cameron & Trivedi IM	No Serial Correlation	SWILK Normality	Overall
CRNZX	HL Sentiment Ratio	F 0.0826	P 1.05	P	P 0.2857	P 0.1543	P 0.693	P 0.60308	F
ARNZXSPY	HL Sentiment Ratio	P 0.4602	P 1.05	P	P 0.7148	P 0.6299	P 0.157	P 0.15056	P
CRTWI	HL Sentiment Ratio	P 0.5393	P 1.05	P	P 0.6461	P 0.7570	P 0.787	P 0.42656	P
CRNZX	HL Subjectivity Ratio	P 0.1611	P 1.02	P	P 0.9434	P 0.8824	P 0.770	P 0.65925	P
ARNZXSPY	HL Subjectivity Ratio	P 0.3433	P 1.02	P	P 0.96917	P 0.6012	P 0.253	P 0.36789	P
CRTWI	HL Subjectivity Ratio	P 0.9250	P 1.02	P	P 0.6668	P 0.7229	P 0.981	P 0.73156	P
CRNZX	HL Polarity Ratio	F 0.0155	P 1.05	P	P 0.1113	F 0.0561	P 0.653	P 0.57805	F
ARNZXSPY	HL Polarity Ratio	P 0.8654	P 1.05	P	P 0.4467	P 0.4112	P 0.145	P 0.14291	P
CRTWI	HL Polarity Ratio	P 0.4873	P 1.05	P	P 0.6809	P 0.8180	P 0.767	P 0.44081	P

Key: P= PASS, F = FAIL

Key: P= PASS, F = FAIL

Appendix D: Correlation Matrix: Zero Conditional Mean

Dependent Variable	Sentiment Variable(s)	Residual to Sentiment	Residual to Press Conference	Residual to OIS OCR
CRNZX	VADER Compound	0.0000 (1.000)	-0.0000 (1.000)	-0.0000 (1.000)
ARNZXSPY	VADER Compound	-0.0000 (1.000)	0.0000 (1.000)	-0.0000 (1.000)
CRTWI	VADER Compound	0.0000 (1.000)	-0.0000 (1.000)	0.0000 (1.000)
CRNZX	VADER Positive	-0.0000 (1.000)	0.0000 (1.000)	0.0000 (1.000)
CRNZX	VADER Negative	-0.0000 (1.000)	0.0000 (1.000)	-0.0000 (1.000)
ARNZXSPY	VADER Positive	0.0000 (1.000)	0.0000 (1.000)	0.0000 (1.000)
ARNZXSPY	VADER Negative	-0.0000 (1.000)	0.0000 (1.000)	0.0000 (1.000)
CRTWI	VADER Positive	0.0000 (1.000)	0.0000 (1.000)	0.0000 (1.000)
CRTWI	VADER Negative	0.0000 (1.000)	0.0000 (1.000)	0.0000 (1.000)
CRNZX	LM Sentiment	-0.0000 (1.000)	0.0000 (1.000)	-0.0000 (1.000)
ARNZXSPY	LM Sentiment	0.0000 (1.000)	-0.0000 (1.000)	-0.0000 (1.000)
CRTWI	LM Sentiment	-0.0000 (1.000)	0.0000 (1.000)	-0.0000 (1.000)
CRNZX	LM Subjectivity	-0.0000 (1.000)	0.0000 (1.000)	0.0000 (1.000)
ARNZXSPY	LM Subjectivity	-0.0000 (1.000)	-0.0000 (1.000)	-0.0000 (1.000)
CRTWI	LM Subjectivity	0.0000 (1.000)	-0.0000 (1.000)	-0.0000 (1.000)
CRNZX	LM Polarity	0.0000 (1.000)	-0.0000 (1.000)	-0.0000 (1.000)
ARNZXSPY	LM Polarity	0.0000 (1.000)	-0.0000 (1.000)	-0.0000 (1.000)
CRTWI	LM Polarity	0.0000 (1.000)	-0.0000 (1.000)	-0.0000 (1.000)
CRNZX	HL Sentiment	0.0000 (1.000)	-0.0000 (1.000)	-0.0000 (1.000)
ARNZXSPY	HL Sentiment	0.0000 (1.000)	0.0000 (1.000)	-0.0000 (1.000)
CRTWI	HL Sentiment	0.0000 (1.000)	-0.0000 (1.000)	0.0000 (1.000)
CRNZX	HL Subjectivity	0.0000 (1.000)	-0.0000 (1.000)	-0.0000 (1.000)
ARNZXSPY	HL Subjectivity	0.0000 (1.000)	-0.0000 (1.000)	-0.0000 (1.000)
CRTWI	HL Subjectivity	0.0000 (1.000)	-0.0000 (1.000)	0.0000 (1.000)
CRNZX	HL Polarity	-0.0000 (1.000)	0.0000 (1.000)	-0.0000 (1.000)
ARNZXSPY	HL Polarity	0.0000 (1.000)	-0.0000 (1.000)	0.0000 (1.000)
CRTWI	HL Polarity	0.0000 (1.000)	-0.0000 (1.000)	-0.0000 (1.000)

Key: Correlation Coefficient (P- Value)

Appendix E: Regression Coefficients and Significance

Dependent Variable	Sentiment Variable(s)	Coefficient	Standard Error	P Value	Statistical Significance
CRNZX	VADER Compound	-0.00221	(0.00174)	>0.1	None
ARNZXSPY	VADER Compound	-0.000158	(0.00274)	>0.1	None
CRTWI	VADER Compound	-0.00196	(0.00206)	>0.1	None
CRNZX	VADER Positive	-0.0305	(0.0339)	>0.1	None
	VADER Negative	0.0149	(0.0468)	>0.1	None
ARNZXSPY	VADER Positive	-0.0349	(0.0522)	>0.1	None
	VADER Negative	-0.140*	(0.0720)	<0.1	* 10%
CRTWI	VADER Positive	-0.0917**	(0.0388)	<0.05	** 5%
	VADER Negative	-0.0716	(0.0535)	>0.1	None
CRNZX	LM Sentiment Ratio	-0.0533	(0.0555)	>0.1	None
ARNZXSPY	LM Sentiment Ratio	0.0847	(0.0866)	>0.1	None
CRTWI	LM Sentiment Ratio	0.138**	(0.0641)	<0.05	** 5%
CRNZX	LM Subjectivity Ratio	-0.000130	(0.000625)	>0.1	None
ARNZXSPY	LM Subjectivity Ratio	-0.00169*	(0.000961)	<0.1	* 10%
CRTWI	LM Subjectivity Ratio	-0.00148**	(0.000720)	<0.05	** 5%
CRNZX	LM Polarity Ratio	-0.00310	(0.00282)	>0.1	None
ARNZXSPY	LM Polarity Ratio	0.00282	(0.00442)	>0.1	None
CRTWI	LM Polarity Ratio	0.00541	(0.00330)	>0.1	None
CRNZX	HL Sentiment Ratio	0.0237	(0.0444)	>0.1	None
ARNZXSPY	HL Sentiment Ratio	0.0928	(0.0688)	>0.1	None
CRTWI	HL Sentiment Ratio	-0.0199	(0.0523)	>0.1	None
CRNZX	HL Subjectivity Ratio	0.000311	(0.000533)	>0.1	None
ARNZXSPY	HL Subjectivity Ratio	-0.000873	(0.000830)	>0.1	None
CRTWI	HL Subjectivity Ratio	-0.00135**	(0.000614)	<0.05	** 5%
CRNZX	HL Polarity Ratio	0.00265	(0.00394)	>0.1	None
ARNZXSPY	HL Polarity Ratio	0.00684	(0.00613)	>0.1	None
CRTWI	HL Polarity Ratio	-0.000589	(0.00465)	>0.1	None

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1